

Sujit Borey, Puja Padiya, Kishor K. Hirekar, Shilpa
D. Chindamwar, Swapnil M. Shelke

Principles of Quantum Machine Learning

Algorithms, Computational Complexity, and Resource Scaling



Principles of Quantum Machine Learning: Algorithms, Computational Complexity, and Resource Scaling

Sujit Borey

ECE department VNIT Ambazari Road 440010 Nagpur

Puja Padiya

Department of Information Technology at Xavier Institute of Engineering, Mumbai

Kishor K. Hirekar

Department of Electrical Engineering, Visvesvaraya National Institute of Technology, Nagpur. Maharashtra

Shilpa D. Chindamwar

Govindrao Wanjari College of Engineering and Technology Nagpur, Maharashtra

Swapnil M. Shelke

Department of Mathematics, Dr. Babasaheb Ambedkar Marathwada University Chhatrapati Sambhajnagar Sub-Campus Dharashiv. Maharashtra



Published, marketed, and distributed by:

Deep Science Publishing, 2026
USA | UK | India | Turkey
Reg. No. MH-33-0658412
www.deepscienceresearch.com
editor@deepscienceresearch.com
WhatsApp: +91 7977171947

ISBN: 978-93-7185-344-6

E-ISBN: 978-93-7185-361-3

<https://doi.org/10.70593/978-93-7185-361-3>

Copyright © Sujit Borey, Puja Padiya, Kishor K. Hirekar, Shilpa D. Chindamwar, Swapnil M. Shelke, 2026.

Citation: Borey, S., Padiya, P., Hirekar, K. K., Chindamwar, S. D., & Shelke, S. M. (2026). *Principles of Quantum Machine Learning: Algorithms, Computational Complexity, and Resource Scaling*. Deep Science Publishing. <https://doi.org/10.70593/978-93-7185-361-3>

This book is published online under a fully open access program and is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0). This open access license allows third parties to copy and redistribute the material in any medium or format, provided that proper attribution is given to the author(s) and the published source. The publishers, authors, and editors are not responsible for errors or omissions, or for any consequences arising from the application of the information presented in this book, and make no warranty, express or implied, regarding the content of this publication. Although the publisher, authors, and editors have made every effort to ensure that the content is not misleading or false, they do not represent or warrant that the information-particularly regarding verification by third parties-has been verified. The publisher is neutral with regard to jurisdictional claims in published maps and institutional affiliations. The authors and publishers have made every effort to contact all copyright holders of the material reproduced in this publication and apologize to anyone we may have been unable to reach. If any copyright material has not been acknowledged, please write to us so we can correct it in a future reprint.

Preface

The intersection of machine learning and quantum computing has been one of the greatest scientific projects in current time. With quantum computing in the offing, the assurance of using quantum mechanical effects in computational learning to transform the world has never been closer. The publication of this book comes at a critical time, when theoretical models are shouldering maturity with capabilities of experimentation, that has brought about great possibilities to implement computational benefits that quantum systems can provide in performing machine learning tasks. Quantum machine learning is one field which has rapidly changed in the last ten years. What was originally viewed as mere theoretical conjecture has now evolved into a living research community which includes rigorous mathematical modeling, novel algorithm design, and more and more complex experimental realisations on the near-term quantum computers. Such change can be characterised by the increasing sophistication of our theory in as well as the revolutionary advances in quantum computing technology, and by the noisy intermediate-scale quantum (NISQ) era.

Chapters in this book chronologically discuss the complex nature of quantum machine learning, including theoretical principles and challenges of practice. All the chapters cover a very important aspect in this growing area leaving the reader not only with an in-depth analysis but also a wide enough point of view to traverse this somewhat complicated inter-disciplinary area. The book is organized in such a way that that it takes the reader through more advanced parts of the quantum machine learning theory and practice. We start by explaining quantum feature maps, and kernel-type methods and laying the mathematical background on which quantum-enhanced learning is based. There we discuss computational complexity, trainability issues, neural network architectures, scaling of resources, sample complexity limits and the issue critical of quantum data encoding.

This book has a variety of audience. The works will also offer exhaustive proposals of key concepts, methodologies and open problems to graduate students and researchers venturing into the field. Individuals who operate quantum computing platforms will have an understanding of the options available when it comes to designing an algorithm and the resources needed, as well as the expectations of its performance. The systematic approach of tackling complexity theorist-foundations and rigorous mathematical analyses will be appreciated by the theorists.

Quantum machine learning is interdisciplinary and requires the services and knowledge of quantum physics, computer science, mathematics, and statistics. Although we do not expect our readers to have the basic knowledge of quantum computing and classical machine learning, we give adequate background and references to accommodate other readers of varying technical background. By 2026, quantum machine learning will be on an important crossroads. NISQ has also provided its own stunning demonstrations and harsh lessons on the difficulty of obtaining useful quantum advantage. Other phenomena like barren plateaus have changed the minds of us on the nature of variational quantum algorithm trainability. Resource-efficient quantum learning has been transformed by the advanced measurement protocols, such as classical shadow tomography. The frontier of the implementable quantum algorithms is constantly being pushed by hardware advances.

The book reflects the picture of the field in this dynamic time, including the latest theoretical developments, their experiments, and new optimal practices. We admit that quantum machine learning is a fast-developing science, and there are certain spheres of our knowledge that will surely become more profound in the nearest future. The transformative potential of quantum machine learning is still achievable only with further achievements in many dimensions: the theory of when and why quantum algorithms are effective, algorithmic innovations to overcome the limitations of near-term hardware, hardware engineering to overcome qubit quality and system scale, and discovering applications when quantum algorithms and methods offer significant superior practical performance.

Hopefully, this book will provide useful guidance to all the existing knowledge and guide to the future innovations. The opportunities are out of the world, yet the challenges are significant. With the ongoing development of quantum hardware capabilities and an ever-increasing theoretical basis, it is quite possible that quantum machine learning will transform the world of approaching the problem of complex computational learning.

Sujit Borey
Puja Padiya
Kishor K. Hirekar
Shilpa D. Chindamwar
Swapnil M. Shelke

Table of Contents

Chapter 1: Quantum Feature Maps and Kernel Expressivity in High-Dimensional Hilbert Spaces1

1 Abstract.....	1
2. Introduction	1
3. The Mathematical Items behind Quantum Feature Maps.	3
4. Kernel Methods and Name Framework Reproducing Kernel Hilbert Space.	5
5. Expressivity Analysis and the Geometry of Quantum Feature Spaces.....	7
6. Quantum Circuit design of Feature map.....	9
7. Measures of Expressivity and Theoretical Limits.....	11
8. Quantum Stimulability Bases vs Classical Advantages.....	14
9. Concentration Phenomena and Barren Plateaus in Quantum Feature Maps.....	17
10. Quantum Kernel Model Training and Quantum Kernel Model Optimization.....	19
11. Generalization Theory and Sample Complexity Bounds.....	21
12. Experimental Demonstrations and Permithive Practical Application.....	22
13. Future Research and Open Research questions.	25
14. Conclusion	27

Chapter 2: Computational Complexity of Quantum Machine Learning Models .29

1 Abstract.....	29
2. Introduction	30
3. Methodology.....	33
3. Results and Discussion	34
3.1 Applications of Quantum Machine Learning Models.....	34
3.2 Techniques in Quantum Machine Learning Complexity Analysis	40
3.3 Methods for Implementing Quantum Machine Learning Models.....	47

3.4 Challenges in Quantum Machine Learning Complexity	53
3.5 Opportunities in Quantum Machine Learning Complexity	58
3.6 Impact of Quantum Machine Learning Complexity Research	63
3.7 Future Directions in Quantum Machine Learning Complexity	68
4. Conclusion	78

Chapter 3: Trainability and Optimization Landscapes of Variational Quantum Algorithms.....81

1 Abstract.....	81
2. Introduction	82
3. Methodology	85
4. Results and Discussion	86
4.1 Theoretical Principles of Trainability in Variational Quantum Algorithms.....	86
4.2 Optimization Landscape Geometry and Characterization.....	90
4.3 Barren Plateau Phenomena and Gradient Vanishing Mechanisms	93
4.4 Mitigation Strategies and Circuit Design Principles	97
4.5 Applications and Domain-Specific Considerations	105
4.6 Hardware Considerations and Noise Effects.....	109
4.7 Emerging Techniques and Future Directions.....	113
5. Conclusion.....	121

Chapter 4: Quantum Neural Network Architectures: Depth, Width, and Expressive Power.....125

1 Abstract.....	125
2. Introduction	126
3. Methodology	127
4. Results and Discussion	128
4.1 Applications of Quantum Neural Network Architectures	128
4.1 Techniques in Quantum Neural Network Architecture Design	131

4.2 Methods for Analyzing and Optimizing Quantum Neural Network Architectures....	134
4.3 Challenges in Quantum Neural Network Architecture Development.....	137
4.4 Opportunities in Quantum Neural Network Architecture Research.....	146
4.5 Impact of Quantum Neural Network Architecture Innovations.....	150
4.6 Future Directions in Quantum Neural Network Architecture Research.....	153
5. Conclusion.....	161

Chapter 5: Resource Scaling in Quantum Kernel Estimation and Measurement Protocols163

1 Abstract.....	163
2. Introduction	163
3. Quantum kernel methods: formalism and measurement primitives.....	164
3. Shot noise, estimator variance, and sample complexity	165
4. Concentration phenomena and implications for scaling	165
5. Measurement protocols beyond single-basis sampling: classical shadows and randomized measurements	167
6. Entangled and joint measurements: trade-offs and advances.....	167
7. Adaptive measurement strategies and estimators	169
8. Error, noise, and mitigation overheads	169
9. Classical postprocessing and computational scaling	170
10. Practical guidelines for resource allocation	170
11. Two comprehensive summary tables.....	171
12. Emerging approaches and research frontiers	173
13. Limitations, open problems, and cautionary perspectives	174
14. Conclusion.....	175

Chapter 6: Sample Complexity and Generalization Bounds in Quantum Learning176

1 Abstract.....	176
2. Introduction	176

3. Preliminaries and formal models	177
4. Key technical tools: quantum analogues of classical complexity measures	178
5. Classical versus quantum sample complexity: equivalences and separations.....	179
6. Representative upper bounds: constructive algorithms and their sample counts.....	180
7. Bases of information theory and adversary dialogue Representation lower bounds.....	182
8. Generalization bounds for parametrized quantum models (QNNs and quantum kernels)....	182
9. The role of noise and NISQ-era constraints.....	183
10. Learning states, learning channels, and learning measurements.....	184
11. Active issues and identification of opportunities.....	184
12. Expansive practice and experiment implications.....	186
13. Conclusion.....	187

Chapter 7: Quantum Data Encoding: Cost and Scalability of Input Preparation
.....188

1 Introduction	188
2. The primary part in end-to-end complexity state preparation.....	188
3. Taxonomy of encoding methods and their resource scalings	190
4. precision versus imprecision in state preparation Trade-offs and algorithms.....	192
5. QRAM architectures, feasibility, and critiques.....	192
6. Compiler and algorithmic techniques lowering costs of inputs.	193
7. Implications of noise, error accumulation and fault-tolerance.....	194
8. Real life examples: machine learning, linear systems and quantum simulations.....	196
9. Conclusion.....	196
References	197

Chapter 1: Quantum Feature Maps and Kernel Expressivity in High-Dimensional Hilbert Spaces

1 Abstract

Quantum computing and machine learning intersect have created quantum machine learning as a rapidly growing area, attempting to apply the physics behind quantum mechanics to improve the learning algorithms. Quantum feature maps and quantum kernels take centre-stage among the foundational constructs in this field as ultimately they have the high-capacity to map classical data onto Hilbert spaces of exponentially large size. These embeddings allow very expressive representations that could be better than the constraints of classical feature spaces and kernel methods. The chapter is a comprehensive and in-depth academic analysis of quantum feature map and kernel expressivity in high dimensional Hilbert spaces with special attention paid to mathematical backgrounds, theoretical expressivity, geometric interpretations, learning-theoretic applications and applications in new places. The role of dimensionality of Hilbert spaces, entanglement, circuit architecture and noise in determining the expressiveness of quantum kernels is especially discussed. Combining both recent theoretical and experimental and algorithmic progress, this chapter provides the general view of opportunities and limitations of quantum-enhanced feature representations and directions on the way to potential practical quantum advantage in machine learning.

2. Introduction

Quantum computing and machine learning at the intersection has become one of the most important opportunities of computational science, and quantum kernel methods are one of the most convinced opportunities of using quantum mechanics principles to recognize patterns and analyze data. The core of this convergence is the idea of quantum feature maps, but which provide the key fundamental mechanism of saturating the exponentially large Hilbert spaces quantum systems are based on with classical data. They use maps over these features which take advantage of quantum mechanics, such as

superposition and entanglement, in order to generate representations of classical computing paradigm inaccessible data. The richness and variety of the functions that can be represented in these quantum embeddings is measured by their expressivity, and this expressivity has become the focus point of study as the quantum machine learning community attempts to find factual quantum benefits in the form of practical applications.

The quantum feature map theoretical framework is based on the mathematical system of quantum mechanics, in which data sample of classical feature spaces are transformed into quantum state in high dimensional Hilbert spaces. This mapping procedure alters the geometrical associations among the data points in such a manner that could probably illustrate patterns and forms that are concealed in standard feature representations. These quantum Hilbert spaces have dimensions that exponentially increase with the size of the qubits in the quantum system whose representational capacity is amazingly well-endowed and far surpasses the representational capacity of the policy feature spaces of classical representations, which are of a size that only grows polynomially with the size of the feature space. Nonetheless, this rapid increase in dimensionality also makes trainability, generalization and practical execution of quantum machine learning algorithms on near-term quantum devices, limited by noise, short coherence times and damp gate fidelities, deeply problematic.

The analysis of the expressivity of kernels in quantum machine learning scenarios has shown that there is a complicated topography of tradeoffs between the theoretical capacity of quantum feature maps and the practical limits of real world quantum devices. Expressivity here is the ability of a quantum kernel to identify the difference between patterns of input, and to estimate the decision surfaces of complex decisions in the feature space. High expressivity is typically a good trait since it allows quantum kernel method to build in more complex relationships between the data, though overfitting can occur with high expressivity in the face of limited training examples being the main cause as compared to the effective de facto dimensionality of the quantum feature space. Moreover, the expressiveness of quantum kernels is closely related to the structure of the quantum circuits to execute the feature maps, and the circuit depth, the types of gates, the behavior of entanglement and parameterization plans all have important implications in understanding the final performance of quantum kernel algorithms.

Recent development of quantum machine learning theory has given more and more sophisticated tools on how to analyze and optimize quantum feature map, based on methods of differential geometry, representation theory, information theory, and statistical learning theory. Such theoretical progress has been coincided by experimental alpha versions of quantum processors on small scale suggesting not only the possibilities but also the weakness of existing quantum hardware in carrying out useful quantum kernel algorithms. Geometric difference A more specific notion of how quantum feature

maps change the metric structure of the input space is considered as a measure that has gained high importance in explaining when quantum kernels can give benefits over their classical counterparts. Equally, concentration studies in the context of quantum Hilbert spaces with high dimensionality have shed light to both the inherent bottlenecks that have to be surmounted in order to obtain the requisite robustness of quantum machine learning performance, whether in the form of barren plateaus or the exponential growth in concentration of the magnitude of quantities being considered.

3. The Mathematical Items behind Quantum Feature Maps.

The mathematical context of quantum feature maps gives the necessary language of how to systematize classical data into quantum states. The essence of quantum feature map is the function that accepts a data point of a classical model, usually specified as a vector of a real or complex vector space, and describes by a collection of quantum states in the Hilbert space of a system of quantum entities. In much stricter terms, assuming we choose to define the classical feature space as a subspace of d dimensional real space, a quantum feature map may be represented as a function that acts between this classical space into the projective Hilbert space of an n -qubit quantum system, having dimension 2^n . The resulting quantum state after this mapping is normally a pure state, which is a normalized Hilbert space or, in other words, it corresponds to a rank-one density matrix.

To define quantum feature maps, the design of parameterized quantum circuits is often the domain of operation, the classical data value values comprise the parameters, and the quantum gates are applied to an initial reference state, often the basis state of the computational basis state, all qubits in the zero state [1-3]. These parameterized quantum circuits comprise unitary transformations which rely on the input data and the geometry and algebraic characteristic of such transformations defines the geometric and algebraic properties of the resultant feature map. The angle encoding technique is a canonical example of a quantum feature map Microsoft Office Microsoft Office is a structural project for remote sensing management that originated more than twenty years ago. An angle encoding encodes classical features as rotation angles of individual qubits using rotation gates, making it a canonical example of a quantum feature map. As an example, what should be the rotation angles in a circuit which uses rotation gates $R_y(x_i)$ or $R_z(x_i)$ on qubits whose components are classical data points (e.g. x_1, x_2, \dots, x_n) i.e. x_n may be used as the angles of rotation.

The action of a quantum feature map is essentially a consequence of the dimensionality and geometry of the image of the classical data of this mapping. In case classical data are encoded as quantum states, the inner products between the quantum states represent a kernel function on the classical feature space, which is called quantum kernel. In more detail, assuming we define the quantum feature map (which maps classical points x to

quantum states) by $|\varphi(x)\rangle$ to be the overlap between quantum states, then the quantum kernel is defined as $k(x, x') = |\langle\varphi(x)|\varphi(x')\rangle|^2$. This quantum kernel has the same properties as the geometry of the quantum feature map, and the geometry of quantum Hilbert space is important in defining the kind of patterns and relations that are possible to learn by this kernel-based technique. The mathematical necessary conditions of a valid kernel functional always positive semi-definite kernels insist on the kernel being positive semi-definite.

Quantum circuit architecture Choosing quantum gates and circuit architecture to compose a quantum feature map has far reaching consequences on expressivity, and computational properties. Single-qubit rotation gates offer a simple model of encoding classical information into quantum states and solely produce product states under isolation. The representational ability of quantum mechanics given its full action requires the addition of entangling gates like controlled-NOT gates, controlled-Z gates, or some more general two-qubit unitaries which introduce correlations between qubits. The entanglements that are created by the quantum circuit create its pattern, which dictates the complexity of the correlations that are possible within the quantum feature map. The quantum kernels of circuits that have a small entanglement can be efficiently simulated, using classical computation devices, which can in turn be a constraint on the quantum advantage available. Alternatively, circuits based on large entanglement can result in very complicated quantum states that are not easily simulated or meanwhile are also more vulnerable to noise and decoherence on current quantum devices.

Methods in the mathematics of quantum feature maps frequently make use of objects in diffusion geometry and Lie theory to describe the manifold of reachable quantum states via circularized quantum circuits, as a parameterized circuit. The tangent plane to this manifold of a specific quantum state gives the information about the local behavior of the feature map and its sensibility to the changes in the input parameters. The main role in this analysis is played by the quantum Fisher information metric that provides a natural Riemannian geometry of the space of quantum states. This index measures the separability of adjacent sized quantum state, and is strongly reliant on the gradient of the quantum kernel function. Generally, feature maps with a large separation in the Fisher information metric make feature maps more expressive because they generate more separation between different input points. Nonetheless, the geometry of the quantum state space is also determined by the geometry of the quantum circuits to realize the feature map and the geometry of schemes of effective quantum machine learning algorithms relies on such geometry constraints.

4. Kernel Methods and Name Framework Reproducing Kernel Hilbert Space.

Kernel methods have a scientific basis that the theoretical framework explains in rigorous mathematical terms the mechanisms of learning tasks in machine learning based on implicit manipulation of high-dimensional feature representations. The kernel methods have also been developed and applied in the classical context and have been effectively tackling various problems in pattern recognition with the support vector machines perhaps the best-known example of the use of a kernel-based learning approach. The kernel methods intuition is that in the vast majority of machine learning algorithms, it is possible to reformulate and entirely rely on inner products between data elements in a feature space, without actually computing the feature representations themselves. This observation is encoded in the representer theorem that demonstrates that optimum solutions to regularized empirical risk minimization are representable in the form of linear combinations of kernel evaluations at the training points.

The kernel trick becomes used differently in the quantum scenario since the quantum feature maps may project data into exponentially larger Hilbert spaces which would be computationally infeasible to code explicitly using a classical memory resource. The quantum kernel function that is defined as the overlap of quantum states created by the feature map can potentially be calculated efficiently on a quantum computer by preparing the quantum states of two data points and measuring the inner product of the two provided by interference. This quantum kernel evaluation needs a number of quantum circuit evaluations which scales exponentially with the precision needed and logarithmically with the dimension of the Hilbert space, which is an exponential improvement over the classical cost of directly computing inner products in a space of exponential dimension. Nevertheless, this possible benefit should be diminished by the awareness that not every quantum kernel automatically makes it better than the well thought-over classical kernels, and what scope of problems quantum kernels really benefit is still in research.

The mathematical environment on which quantum kernel methods are applied is the reproducing kernel Hilbert space of a quantum kernel. This Hilbert space is comprised of functions that take as input the classical space, to real or complex values, with inner product, defined by the kernel function. The reproducing property has the effect that assessment of functions in this space at a given point can be expressed as an inner product with the kernel function assessed at the point. In the case of quantum kernels, functions in the RKHS may be a linear combination of evaluations of a kernel, and the strength of such a combination is the measure of the expressivity of the quantum kernel method. The dimension of the RKHS has to do with the rank of the kernel matrix on the data set, and in the case of quantum kernels being built out of feature maps into high-dimensional quantum Hilbert spaces, this dimension can be huge, and so may enable the expression of highly complicated decision functions.

The relationship between quantum kernels and classical kernels can be mathematically operationalised (using) the concept of kernel alignment and kernel similarity measures. Quantum kernels that are closely related to efficient classical kernels, including those based on polynomials, or Gaussian kernels, are hardly likely to offer computational benefits due to the ability to compute classical kernels with high efficiency without using quantum computation. Quantum advantage in kernel methods is thus pursued by the construction of quantum feature maps that have properties which generate quantum kernels with unique properties which are difficult to easily approximate using classical kernel functions. This results in taking the expressive power of quantum kernels this is aimed at characterizing the kinds of functions and decision boundaries which are well represented in the RKHS of quantum kernels but not represented in the RKHS of classical efficient kernels. Instances of such functions, and their application to real-world machine learning tasks, is a classical problem in the quantum machine learning methodology.

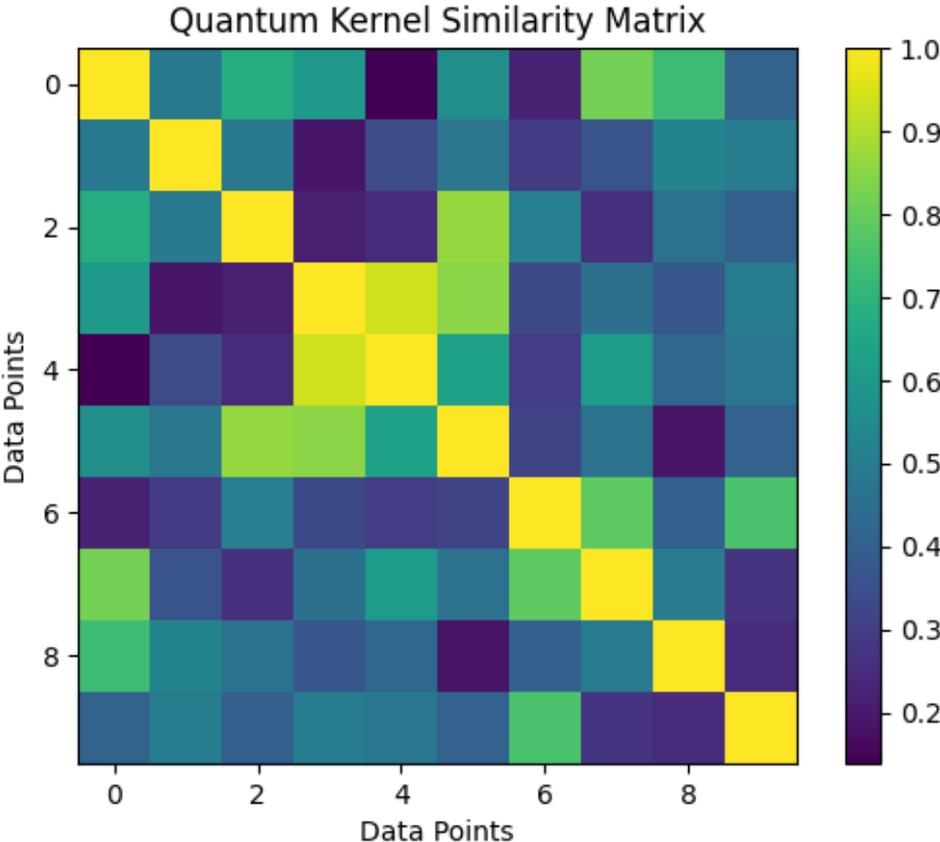


Fig 1: Pairwise Kernel Similarity Heatmap (Quantum Kernel Matrix)

The complexity of computing quantum kernels on quantum devices is built upon the circuitry of quantum features map. On product state feature map S mapping classical

data to quantum states, the product state quantum kernel may be efficient to compute (with respect to feasible classical computation) and preclude any possible quantum advantage. In the case of feature maps that create entangled states the quantum kernel evaluation can demand genuinely quantum organizational resources however, the entanglement itself does not imply computational benefit or other performance. The size of the quantum circuits, the number of the qubits used, and the type of a gate set used all affect both theoretical and practical expressivity of the quantum kernel and its practical feasibility on near-term quantum implementations. More noise-resistant shallow quantum circuits with smaller entanglement can be easier to implement, although can also define quantum kernels with limited expressiveness that find limited application over the classical approach.

5. Expressivity Analysis and the Geometry of Quantum Feature Spaces

The expressibility of quantum feature maps and feature kernel itself is a complicated notion that implies a number of different yet similar notions of representational capacity and discriminative power. On the simplest level, expressivity can be seen as the variety of the quantum states that can be obtained as the classical input of a quantum feature map ranges through its space. The rich repertoire of quantum states accessible to a highly expressive feature map is well-distributed across the accessible range of the quantum Hilbert space, and allows the quantum kernel to detect more complicated patterns and relationships in the data. On the other hand a reduced expressivity feature map can also result in quantum states which are restricted to a low dimensional submanifold of the Hilbert space, potentially limiting the capacity of quantum kernel method to differentiate between varying input patterns or model complex decision boundaries.

A strict method of measuring expressivity is to compute the rank and spectrum of the quantum kernel matrix which is the Gram matrix of inner products of quantum states representing each pair of data points of a given dataset. This kernel matrix is an $n \times m$ Hermitian positive semi-definite matrix (m points), whose eigenvalues and eigenvectors carry significant information as to the geometry of the quantum feature space. The rank of the kernel matrix, also called the count of non-zero eigenvalues will give an indication of the effective dimensionality of the space that is spanned by the quantum feature representations. When a complete kernel matrix has been filled, then it means that every point in the data has been represented by a linearly independent set of quantum states which points to high expressivity and when the kernel matrix has low rank it means that the quantum states are confined to a subspace of low dimension, indicating low expressivity. Nevertheless, this rank measure is to be understood with caution since the quantum hardware may produce noises, as well as limited precision,

which directly results in the approximate collapse of eigenvalues, which may cause artificial inflation of the numerical rank.

Another significant view of the expressivity is the geometric organization of the space of quantum states reachable with the help of a parameterized quantum circuit. Such a quantum state manifold is a part of the projective Hilbert space and its dimension, curvature, and topological properties will specify the kind of quantum states that can be generated as well as the association between them. The quantum state ball of a quantum circuit, which has a number of independent parameters p , is generically a p -dimensional submanifold of the projective Hilbert space which is itself fiberwise the real dimension $2^{n+1} - 2$ of a n -qubit system. The rate of the ratio between the dimension of the accessible manifold and dimension of ambient Hilbert space gives one measure of the coverage of the Hilbert space by the feature map. Quantum circuits deep enough and with suitable gates structures may be able to cover the Hilbert space well, giving out quantum state manifolds that span a high volume of the accessible quantum state space.

The idea of expressivity can be associated with the idea of universality in quantum computing, which is the fact that a given set of gates can serve to approximate any unitary transformation on the quantum Hilbert space to any desired degree of precision. Universal gate sets make it possible to create quantum circuits, which can create any quantum state, up to a total phase, by implementing an appropriately long sequence of gates. Nonetheless, the notion of universality here is not the same as expressivity of quantum feature map since the latter is about the particular mapping of classical data to quantum states but not about the ability to create arbitrary quantum states. A quantum feature map may have a universal set of gates, but still limited expressivity when the scheme of parameterization does not be able to explicitly control the parameters of a quantum circuit such that the classical input data can ensure a wide range of quantum states.

Expressivity analysis also needs to take into account those phenomena of concentration of measure that take place in high dimension Hilbert spaces of quantum. Random quantum states with higher and higher qubits are closer and closer together in their pairwise overlaps, with most of the overlaps clustering around zero. This concentration phenomenon may impact the quantum kernels such that the values of the kernel become concentrated around a fixed value which narrows the capabilities of the kernel to distinguish between various input points. Quantum states with the concentration behavior of feature maps can be poorly expressed in practice even though their space of particularly simple states occupies a high-dimensional manifold of the Hilbert space. The design of quantum feature maps to reduce the effects of concentration should be done with care, since the quantum states of different classical inputs need to be large enough to be distinguished.

6. Quantum Circuit design of Feature map.

Quantum feature maps in measurement In order to put quantum feature maps into practice, quantum circuits must be designed, which are able to code classical information in quantum states and have enough expressivity and resilience to hardware noise [2,4]. There are a number of important design choices in the architecture of such quantum circuits, such as the choice of encoding gates, entanglement and qubit connectivity pattern, the circuit depth and layering structure, and a parameterization strategy. All these decisions have an effect on the theoretical properties of the resulting feature map, as well as on its practical implementation on the near-term quantum devices. The conflict between expressiveness and the need to ensure the remaining implementability in noisy intermediate-scale quantum devices is also a primary issue in the machine learning context of quantum circuit design.

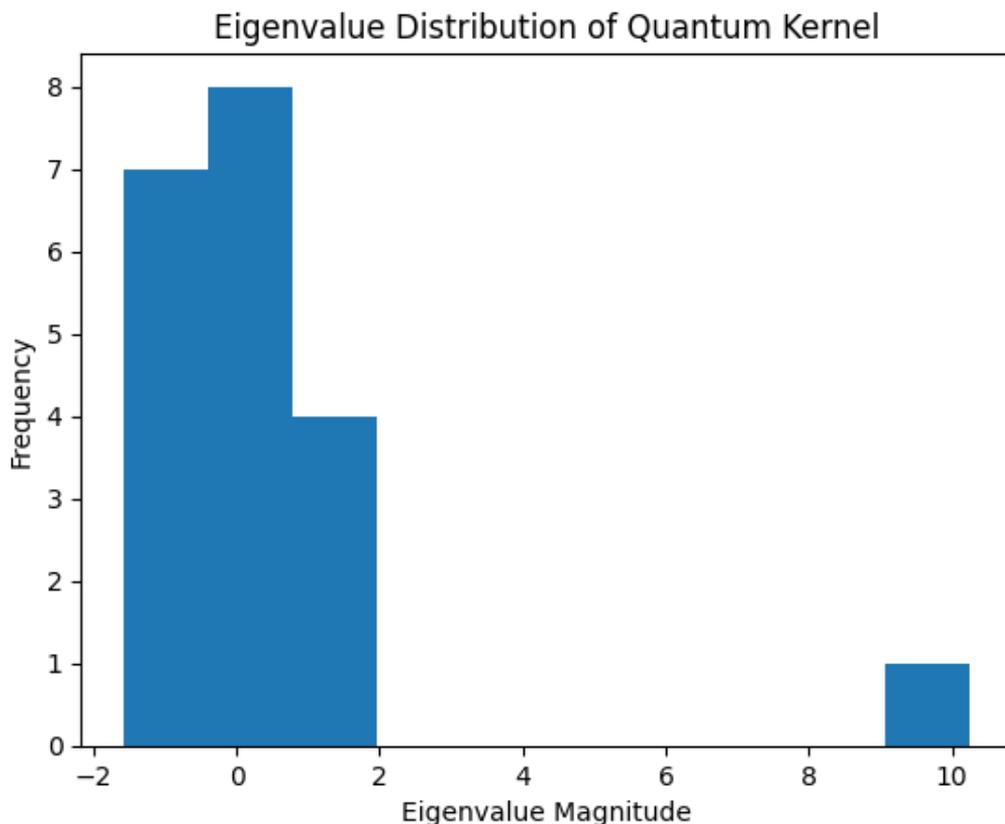


Fig 2: Distribution of Kernel Eigenvalues (Expressivity Analysis)

Quantum feature map circuit design is based on data encoding strategies, which define the way classical information is first encoded onto the quantum system. The simplest method is the basis encoding, in which classical bit strings are directly encoded as computational basis states of the quantum system but this encoding is only appropriate

on discrete data sets and not optimally makes use of the continuous character of the quantum amplitudes. Amplitude encoding maps data in a classical parameterized computation to the amplitudes of a quantum state, and exponentially many classical parameters can be represented in a quantum state of logarithmic numbers of qubits, however, the encoding computations, which use complex state preparation methods, can be difficult to execute efficiently. A more viable middle ground that is commonly used in quantum machine learning applications is angle encoding, in which classical features define the rotation angles of quantum gates. This encoding permits the capability of, in a continuous classical setting, encoding a sequence of qubits as points on the Bloch sphere and the sequence and rotation axes can be selected to maximize the properties of the resulting feature maps.

Entangling operations Adding entangling operations to quantum feature map circuits are necessary to have expressivity that cannot be achieved by classical kernel methods. Entanglement forms correlations between qubits which are inseparable into independent single-qubit states thus allowing the quantum feature map to derive complex interactions between various features of the classical data. The literature has explored many different entangling gate patterns, including the simplest cases of nearest-neighbor entangling interaction structures, or more complex structures that can be in the form of all-to-all entanglement with every pair of qubits. Linear entanglements, in which qubits are configured in a daisy chain, and entangling gates are interacting between adjacent qubits, offer a hardware-efficient implementation that can be accommodated by the small connectivity of many existing quantum devices. **Linear Patterns** Circular patterns of entanglements follow the linear one by linking the first qubit with the last one to provide more ways in which quantum information can be conveyed through the circuit. Patterns of star entanglements, in which a single qubit is entangled with all other qubits, are capable of correlational distribution efficiently, but can effectively form a bottleneck at the central qubit.

Layering and circuit depth are important attributes of defining the expressivity as well as noise robustness of quantum feature maps. Multi-layer more circuit depth gate-based feature maps could in theory be more expressive by enabling the quantum state to be subjected to more complicated transformations as it passes through the circuit. Nevertheless, more complex circuits obtain a greater amount of noise and decoherence, especially on contemporary quantum devices where gate fidelities are not perfectly defined and coherence times are short. The trade-off between these competing considerations is the optimal circuit depth, which varies depending on the details of both the medical device and the machine learning application. **The circuit architecture** Layered circuit architecture: In a layered circuit architecture, the circuit is represented as a sequence of repeated blocks of gates with a regular structure, which is systematically built to represent a quantum feature map of controlled depth. All the layers usually have

a parameterized rotation block including single-qubit rotation using the classical input data, and then entangling block generating correlations among qubits.

The strategies of parameterization specify the way in which the classical input data is used to parameterize the quantum gates in the feature map circuit, and the various ways of parameterizing may give qualitatively different properties of feature maps. Direct parameterization represents the classical features as rotation angles directly in the quantum gates which form a simple map between classical data and quantum states. The method is easy to understand and interpret however, it is not the best method of bringing out high representational power of the quantum system. Learned parameterization wants to add more trainable parameters to the quantum circuit which gain the benefit of being optimized over the course of the training program and it enables the feature map to get acquainted with the particular structure of the data. Nevertheless, this method does not separate the roles of design feature maps and model training, and in cases with small training data that can expose it to overfitting. Hybrid setting Hybrid parameterization methods involve both direct encoding of classical features and trainable parameters with multiple methods trying to trade off the best of both worlds.

7. Measures of Expressivity and Theoretical Limits.

To measure the expressiveness of quantum kernels, it is necessary to devise mathematical measures to obtain the appropriate values of such measures and how these values relate to the structure of the underlying quantum feature maps. The quantum machine learning literature has suggested many different expressivity measures, each targeting one of the aspects of the representational power of the kernel. The effective dimension of the kernel is one such measure that is associated with the trace of the kernel matrix of Frobenius norm which measures the number of effective degrees of freedom of the kernel-based model. An effective dimension is high, which means the one is more expressive i.e. the kernel has the capacity to pick out more patterns within the data. Nevertheless, too much effective dimension in comparison to the amount of training examples may cause overfitting when the model is overfitted i.e. too much like the training examples and thus is unable to predict the new unseen data.

The spectral characteristics of the quantum kernel matrix give finer details of the expressivity in terms of the inequality distribution of the eigenvalues. A slow rate of decaying eigenvalues in a kernel matrix, a large number of which are much greater than the subsequent ones, is a sign of deep dispersion of information across many dimensions of the feature space and so high expressivity. On the other hand, having a kernel matrix whose eigenvalues decay rapidly so that most of the information is concentrated in few leading eigenvectors shows appearance of lower effective expressivity in spite of the possibly large nominal dimensionality. The eigenvalue spectra can be analyzed to

determine whether a quantum feature map is making use of the full feature of the quantum Hilbert space or the quantum states generated by the feature map are limited to a low-dimensional subspace. Ongoing theoretical efforts have obtained constraints on the spectral of eigenvalues of quantum kernels using the geometry of quantum circuits underlying the feature map, which have proven useful in prediction of quantum circuit architecture of kernel expressivity.

Geometric expressivity measures are concerned with the transformation of quantum feature maps of the input space with respect to the metric of the input space, in terms of embedding the input space in quantum Hilbert space. Quantum kernel is a defined distance measure of the classical input space by relation between values of kernel and distance of Hilbert space. In particular, the distance between quantum states, squared, could be represented by the expression two minus twice the real part of the amount of the kernel. It is preferable that feature maps that alter the distance relationship of points relative to one another significantly when compared to the distance relationships in the original classical space be considered more expressive because they define new geometric structures that can be more compatible with the learning task. Geometric difference, which is used to measure the distance between the classical distance measure and quantum kernel-induced measure has become a useful technique to understand when quantum kernels perform better than classical measures. Large geometric difference This indicates that it is a quantum feature that is revealing patterns which would not be seen in the classical feature space.

Information-theoretic measures of expressivity are inspired by quantum information theory ideas of the description of the ability of quantum feature maps to encode and to process information. One measure of this is the von Neumann entropy of quantum states generated by the feature map, and the larger the entropy the more typically the state of the system is mixed and the more complex. As deterministic feature map generation often produces pure states, von Neumann entropy equals zero in this case, but the entropy of reduced density matrices resulting by tracing out jobs of qubits may be non-zero, and will provide information about entanglement structure. Another expressivity measure in information theory of the feature map is the mutual information of the various partitions of the quantum system where it corresponds to the strength of the correlations formed between the feature map and the original system. The ability of feature maps to represent complex functional relationships may be enhanced by feature maps forming large scale entanglement and correlating with each other.

Theoretical limits on quantum kernel expressivity have been considered founded on the underlying limitations of quantum mechanics and circuit structure of quantum circuits. The size of the quantum Hilbert space gives a definite maximum limit to the rank of the quantum kernel matrix, although it may be further decreased based on the utilization of the feature map circuits structure. In case the number of parameters in a quantum circuit

is fixed, the number of parameters in the manifold of reachable quantum states is no longer than the number of parameters, which is a constraint upon the expressivity of quantum circuits on the Hilbert space of any size. Inequality of concentration gives probabilistic upper bounds on the amount of difference in how the value of a kernel spreads around its value in the mean, and it has consequences on distinguishability of various inputs. Recent efforts also had found some basic trade-offs between various desirable quantum kernel properties where maximizing certain qualities of quantum kernels could trade off other qualities such as being trainable on or performing generalization.

Table 1: Comparative Analysis of Quantum Feature Map Architectures

Architecture Type	Classical Simulability	Primary Applications	Computational Complexity	Implementation Challenges	Theoretical Advantages	Notable Limitations
Amplitude Encoding	Exponentially hard for large n	High-dimensional data compression	$O(\text{poly}(n))$ preparation	Complex state preparation	Exponential feature space	Difficult preparation
Basis Encoding	Trivially classical	Discrete categorical data	$O(n)$ for n features	Limited to discrete data	Simple implementation	No continuous encoding
Angle Encoding	Classically tractable	Feature engineering tasks	$O(nd)$ gates	Moderate gate overhead	Hardware efficient	Limited entanglement
Linear Angle Encoding	Hard for deep circuits	Pattern recognition	$O(nd)$ gates	Requires wraparound connections	Enhanced connectivity	Connectivity constraints
Circular						
IQP-Style Encoding	Conjectured hard classically	Quantum advantage tasks	$O(n^2d)$ two-qubit gates	Requires many CZ gates	Theoretical hardness	Susceptible to noise
Hardware Efficient Ansatz	Problem dependent	Near-term applications	Hardware dependent	Training optimization	Tailored to devices	Hardware specific
Quantum Random Kitchen Sinks	Classically hard for large systems	Kernel approximation	$O(\text{poly}(n, m))$ for m measurements	Requires multiple measurements	Theoretical guarantees	Measurement overhead
Separable Feature Maps	Efficiently classical	Baseline comparisons	$O(n)$ single-qubit gates	Minimal but limited power	Easy to implement	No quantum advantage

Tensor Network Inspired	Hard for high bond dimension	Structured data	Depends on contraction	Complex design process	Exploits data structure	Design complexity
Data Re-uploading	Generally hard	Universal approximation	$O(nLd)$ gates	Many parameterized layers	Universal representation	Deep circuit requirements
Quantum Convolutional	Hard for full circuits	Spatial data processing	$O(n \log n)$ typical	Quantum pooling design	Exploits locality	Limited to structured data
Hamiltonian Evolution	Hard for many-body systems	Quantum simulation tasks	Depends on Trotter steps	Simulation accuracy	Physical motivation	Trotter errors

8. Quantum Stimulability Bases vs Classical Advantages.

The issue of whether quantum kernels can be beneficial compared to classical kernel methods is one of the most far-reaching and most debated problems in the research of quantum machine learning. To come up with a rigorous quantum advantage, one would have to prove that quantum kernel methods can solve some form of learning problems at least significantly better than any classical algorithm perhaps in terms of computational resources or in terms of the quality of the learned model. But establishing such benefits has turned out to be quite a challenge and the architectural landscape of quantum kernel algorithms is one where there is a multifaceted relationship between what is not feasible in theory and what is infeasible in practice. This is complicated by the point that quantum computers are not universally better than the classical ones, but can be particularly effective at some types of problems, which can take advantage of the peculiarities of the quantum mechanics.

Quantum stimulability Classical stimulability of quantum kernels is a crucial parameter in determining quantum advantage because quantum kernels that can be efficiently calculated by classical algorithms do not benefit in any way in terms of their quantum hardware. Various noteworthy findings have put in place a set of conditions in which quantum kernels are still classically stimulable, effectively highlighting parts of the quantum circuit design space in which quantum advantage will never exist. Quantum feature maps to fix the problem of entangling qubits, and that would create separable states, are able to generate kernels that are representable as a product of one-qubit terms, and thus can be efficiently evaluated (classically). There are also some classes of shallow entangled quantum circuit of which kernels can be given efficient classical representations with assistance of a tensor network or other means. These classical

stimulability findings indicate the need to consider the depth of circuits as well as the right entanglement structure towards attaining potential quantum benefits.

A common belief on this subject, that a quantum kernel, which is difficult to compute in classical computations but is still useful in machine learning, is a fine balance that is not necessarily satisfied by arbitrary quantum circuits. Computational hardness by itself does not imply that quantum kernel will be more useful in actual machine learning problems since features induced by the kernel will be useful only to the extent that they are relevant to the problem under solution. This realization has created the study of quantum kernel alignment which looks at the interaction between quantum features maps and the target functions or decision bounds that must be acquired. Highly alignment kernels to the learning task have higher chances to show good performance, but the high alignment has to be obtained and at the same time computational hardness with comparisons to classical methods to provide a quantum advantage.

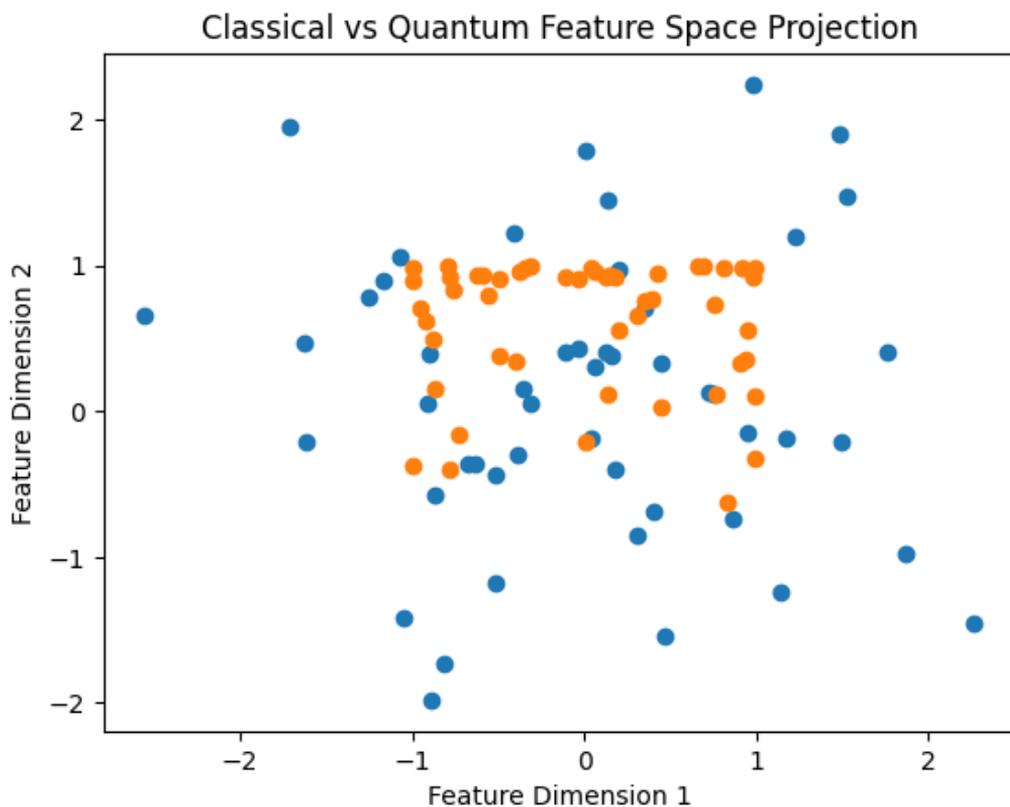


Fig 3 : Classical vs Quantum Feature Space Separation (2D Projection)

Computational complexity theory has been used to develop theoretical frameworks of the quantum advantage of kernel methods, especially of sampling problems and oracle separations. There is an influential method that forms a learning problem by data being

generated by quantum processes considered difficult to model classically, This reply stopped since Claude was at its maximum length of a message. Hit continue to push Claude on. Continue Claude is a computer that can be at fault. Checking Responses: Please check the responses.

Table 2: Expressivity Metrics and Performance Characteristics of Quantum Kernels

Expressivity Metric	Mathematical Definition	Relationship to Performance	Interpretation	Theoretical Bounds	Optimization Strategy	Practical Considerations	Measurement Requirements
Kernel Matrix Rank	Number of non-zero eigenvalues $\lambda_i > \epsilon$	Higher rank enables complex functions	Effective dimensionality	Bounded by $\min(m, \dim(H))$	Maximize while avoiding overfitting	Numerical precision critical	Requires full kernel matrix
Effective Dimension	$\text{tr}(K)^2 / \text{tr}(K^2)$	Optimal intermediate values	Degrees of freedom	Bounded by matrix rank	Balance with sample size	Trade-off with generalization	Computable from kernel
Geometric Difference	Avg divergence of classical and quantum distances	Larger values suggest advantage	Metric transformation strength	No tight general bounds	Maximize for novel patterns	Problem dependent meaning	Distance in both spaces
Average Kernel Value	Mean of $k(x_i, x_j)$ over all pairs	Affects concentration behavior	Overall state distinguishability	Approaches 2^{-n} for random states	Avoid extreme concentration	Connects to barren plateaus	Quantum circuit measurements
Kernel Alignment	$\langle K, Y \rangle / (\ K\ \ Y\)$ for label matrix Y	Higher alignment improves learning	Correlation with target	Maximized at perfect match	Maximize for given task	Task-specific metric	Requires labeled data
Spectral Entropy	$-\sum_i (\lambda_i / \text{tr}(K)) \log(\lambda_i / \text{tr}(K))$	High values suggest expressivity	Information capacity	Maximized for uniform spectrum	Encourage diversity	Dimensionless measure	Full spectrum needed
Fisher Information	$\text{Tr}(\rho^{-1} (\partial p / \partial \theta)^2)$ averaged	Large values aid optimization	Parameter sensitivity	Quantum Cramér-Rao bound	Avoid vanishing gradient	Connected to trainability	Gradient computations

Expressibility	1 - fidelity between ensemble and Haar measure	Higher suggests universality	Hilbert space coverage	Unity for perfect coverage	Increase with depth	Approximates t-designs	Many circuit instances
Entanglement Capability	Average Meyer-Wallach measure	Correlates with non-classicality	Degree of entanglement	Unity for maximally entangled	Maximize for advantage	Exponential in qubit number	Tomography or witnesses
Concentration Coefficient	Variance of kernel values / mean ²	Low values problematic	Kernel value spread	Approaches zero exponentially	Avoid extreme concentration	Indicates barren plateaus	Statistical from kernel values
Generalization Gap	Training accuracy - test accuracy	Lower gap indicates better generalization	Overfitting measure	Statistical learning bounds	Minimize via regularization	Fundamental to ML success	Separate train/test sets
Projected Quantum Dimension	Dimension of span of quantum states	Higher allows richer functions	Accessible Hilbert space	Bound by total dimension	Maximize accessible space	Exponentially hard to measure	Quantum state tomography

9. Concentration Phenomena and Barren Plateaus in Quantum Feature Maps

In the focus of studying the phenomena of concentration in high-dimensional quantum Hilbert space it has been found out that there are inherent challenges that have to be challenged when implementing quantum machine learning models, as feature maps. The dimensionality of the Hilbert space of a quantum system grows exponentially with the number of qubits, and these geometric and statistical characteristics are counterintuitive which can have a critical negative effect on quantum kernel methods. Among the most important concentration effects is the distribution of pair-wise overlaps of the quantum states which are randomly chosen in the Hilbert space. In large quantum systems, these overlaps plunge sharply around their mean value which approaches zero with an increase in the number of qubits, and exponentially small variance. This concentration means that randomly selected quantum states are almost orthogonal to each other and that a perturbation of the quantum state to a small amount has insignificant effects on the overlap of quantum states.

Google The concept of kernel value concentration where quantum kernel evaluations can be concentrated around a narrow range of mean values despite the nature of the input data is a particular form of concentration of measure which has a direct effect on machine learning performance. In case a quantum kernel is strongly concentrated, it means that the numbers of kernels of various pairs of data points are almost the same, which destroys the ability of the kernel to distinguish between various inputs and makes it useless in the learning process. This concentration may occur through a combination of different factors, which include the architecture of a quantum circuit, the size of the quantum system compared to the volume of data and the statistical characteristics of the data distribution. The case of feature maps producing roughly randomly distributed quantum states which are uniformly distributed in the Hilbert space is particularly prone to the effects of concentration, because the random state ensemble has the concentration effects of high-dimensional spaces in it.

The closely related phenomenon is that of barren plateaus, which appears when parameterized quantum circuits are being trained with an objective function that has gradients that go to zero exponentially higher with the number of qubits. Gradient gradient cancellation eliminates the ability of gradient-based optimization to scale to good parameter values exponentially, and in any event makes the process of training deep quantum circuits on large quantum systems prohibitively expensive. Although barren plateaus were previously studied as variational quantum algorithms with trainable parameters, they are also extended to quantum feature maps consisting of both data-dependent and trainable parameters. When a quantum feature map is modeled in a way that the gradient of the parameters has barren plateau behavior, then it is effectively impossible to optimize the parameters to improve the output of the feature map, which is a serious inhibitor to the viability of adaptive or learned quantum feature maps.

Expressivity and concentration do not go hand-in-hand, which is a principal diad at quantum feature map design. More expressive quantum circuits capable of reaching a large fraction of Hilbert space and be able to generate a variety of quantum states are also more susceptible to concentration effects and barren plateaus in that their expressivity may be based on the construction of approximately random quantum states. On the other hand, quantum circuits with low expressivity (that is, induce states in a low-dimensional submanifold of the Hilbert space) can escape concentration and will not be sufficient to learn complex tasks. This poses an optimization problem on the quantum circuit architecture which needs to be optimally configured to obtain enough expressivity to address the learning task, but also not to fall into the pathological concentration behavior that would render the resulting quantum kernel ineffective.

The recent theoretical and empirical studies have pinpointed a number of strategies relating to reducing the effect of concentration and barren plateau in quantum feature maps. By making sure that local parameters do not disappear as the global parameters

do, it is possible with the help of local cost functions which are defined only in terms of measurements of small sub systems and not global ones. Properly designed quantum circuit architectures, including those with finite range of entanglements or with hierarchical or hierarchical structures are able to circumvent the conditions that result in excessive concentration and still express themselves usefully. The strategy of initializing quantum circuits with quantum circuits that are nearer to classical algorithms or with problem structure can be useful to explore the parameter space better. Also, concentration can be alleviated by using classical pre-processing to simplify the resulting effective dimensionality of the input data prior to quantum encoding; this will actually simplify the learning task the quantum feature map will have to complete.

10. Quantum Kernel Model Training and Quantum Kernel Model Optimization.

The machine learning model training and optimization of quantum kernel models can be considered to have a set of different yet connected challenges that include efficient evaluation of the quantum kernels to solving the resultant classical optimization problems. In contrast to variational quantum algorithms which use parameterized quantum circuits, the parameters of quantum circuits are optimized during training, quantum kernel methods use fixed feature maps whose parameters are usually selected by the data directly (or by a meta-optimization process). Solving a convex optimization problem in the reproducing kernel Hilbert space is the main training problem in quantum kernel methods which may be expressed either as support vector machine training, kernel ridge regression or other forms of learning based on kernels. Though such classical optimization problems are known to be well understood and efficient algorithms can be used to solve them, computation using quantum kernels raises new computational concerns such as the cost of evaluation of kernel and the structure of the kernel matrix.

The calculation procedure of training quantum kernel models starts with the calculation of the quantum kernel matrix, which involves the calculation of the values of the kernel between every pair of training samples. Every kernel evaluation is associated with the preparation of quantum states one corresponding to two data points, running the quantum feature map circuits, and measuring the overlap of the resultant quantum states by way of an interference measure, or swap test. In a training set of m examples, this takes $O(m^2)$ quantum kernel evaluations, each requiring the execution of quantum circuits on quantum hardware. The time to execute a quantum circuit is $O(m^2 \times \text{circuit depth} \times \text{shots/eval})$, and the quantum circuit precision of each kernel estimate is determined by the number of shots. This quadratic growth in the size of the training set sometimes can be a bottleneck to large datasets, so research on approximate kernel methods and sampling techniques has been investigated to minimize the number of kernel evaluations needed.

After computing the quantum kernel matrix, the training problem can be solved as a classical optimization problem that can be solved with the help of common machine learning methods. In the case of support vector machines with quantum kernels training consists of solving a quadratic programming problem that optimally solves the linear separator of the quantum feature space, by subject to margin maximization as well as constraints on the slack variables. Countless of such classical optimization problems can be efficiently dealt with as sequential minimal optimization, interior point computing techniques or special-purpose SVM solvers, and the computational complexity of exact algorithms might be $O(m^3)$ and of approximate algorithms might be $O(m^2)$ to $O(m^{2.5})$. These classical optimizers are fed with the quantum kernel matrix, the properties of which including conditioning and spectrum depend on the convergence and numerical stability of the optimization process. Poorly-conditioned quantum kernel matrices whose eigenvalues are too big or too small may cause numerical problems in the training optimization problem.

Hyperparameter optimization is another significant challenge of training quantum kernel models, including the regularization parameters, parameters of the kernel when its feature map can be tuned (e.g., circuit depth of entanglement structure), as well as the architecture. The regularization parameter that determines the compromise between modelling the training data and keeping models simple is especially important to quantum kernels because of the dimensionality of the quantum feature space and oversizing. Cross-validation offers a standardized method of hyperparameter selection, with the training data being divided into folds and model performance is measured on withheld validation folds using selection of varying hyperparameter values. The computational cost of cross-validation is however multiplied by the configurations of hyperparameters that are searched and since quantum kernel evaluations are expensive, this may result in ban some overall training times.

Quantum kernel model Advanced training methods have also investigated different ways of enhancing efficiency and performance. Quantum kernel alignment methods are aimed locally to find the optimal quantum feature map parameters, to maximize quantum aligning with the ideal kernel to the learning task, as characterized by the label information. This meta-learning problem can be described to be an optimization problem at the bi-level in which the outer level is the optimization of the feature map parameters and the inner level is the optimization of the kernel learning problem with the feature map parameters fixed. The other direction is the creation of online and incremental learning algorithms of quantum kernels, which update the model on receiving new training examples instead of retraining and remodelling it at that point. These algorithms are capable of reducing the computational cost of training, by skipping the rapid re-computation of the kernel values between previously trained examples, at the cost of potentially reduced optimality to batch methods of training.

11. Generalization Theory and Sample Complexity Bounds

Quantum kernel methods generalization performance, a measure of the method making correct predictions on unseen data, having the ability to predict with the aim of a finite sample, is controlled by the basic principles of the statistical learning theory that needs to be modified to reflect the specifics of quantum feature maps. The classical findings of operation of kernel methods are a starting point of this analysis, with the exponentially large Hilbert spaces and unique geometries of quantum kernels bringing new issues. The generalization error, the loss that is expected to occur on the data distribution difference between the training loss, is sensitive to a number of major factors such as the complexity of the hypothesis class induced by the quantum kernel, the amount of training data, the noisiness of the data and the quantum hardware as well as noisiness properties of the learning algorithm trained.

Sample complexity bounds have been used to give quantitative approximations of the number of training examples needed to reach a given level of generalization performance and these bounds are needed to have an idea about the practical usefulness of quantum kernel methods. In the classical model, the complexity of a sample in relation to kernel methods is usually described by such concepts as the Rademacher complexity, the covering numbers of the function class, which is the effective capacity of the hypothesis space. In quantum kernels, the large dimensionality of the quantum Hilbert space may be taken to indicate enormous complexity, and a large sample noise. But the effective dimensionality of the upwardly accessible quantum kernel dependent on the nominal dimensionality of the Hilbert space is not the dimensionality of the functional accessible with that quantum kernel, but instead the quantity of the kernel alignment or spectral properties of the kernel matrix or the distribution of margins in classification problems.

Fully marginal theory of generalization This theory has been highly successful in describing performance on classical support vector machines, it generalizes to the quantum kernel methods. The distance between the example and the decision boundary in the feature space is the margin of a training example, and large margin classifiers which separate the training data by large margins are well generalized compared to classifiers with small margins. In the case of quantum kernels margin values are quantified on quantum Hilbert space geometry of the quantum feature map and the distribution of margin values on the training set will inform one regarding the overall generalization capacity. It can be seen that margin-based generalization bounds have affirmed that generalization error reduces with an increase of minimum margin and reduction of effective dimensionality, which is defined by the trace of a kernel matrix, and sample size. These bounds indicate that quantum feature maps that produce well separated quantum states with large pairwise distances are also theoretically capable of having good generalization despite having limited training data.

The issue of overfitting is a major issue of quantum kernel methods since there is a potentially enormous scale of quantum feature spaces. Through the effective dimensionality of the quantum kernel, the model may induce CAPM when the quantum kernel of the observed training examples and when the dimensionality of the quantum computing device are on par with or above the number of training examples, resulting in the model correctly learning the training data but no longer being able to manage the underlying pattern of appearance and behavior of the observed world. Regularization methods offer the main tool of regulating the problem of overfitting with kernel methods and the role of regularization of the overfitting parameter regulates the compromise between the performance of the training and the model complexity. In the case of quantum kernels, the regularization parameter should take both into consideration the nominal capacity of the quantum Hilbert space, and the effective capacity that is given by the particular quantum feature map and data distribution. According to the theoretical analysis, the more effective dimension the quantum kernel is or the fewer training data there are, the deeper regularization is required.

Another issue that provides even more complexity to the theoretical discussions of quantum kernel methods is the effect of quantum hardware noise on the generalization performance. Noise in quantum circuits causes an error between the ideal quantum states which would be achieved by an ideal quantum computer and the noisy quantum states achieved using actual hardware. These deficiencies are flaws in the evaluations of the kernel, in which the measured values of the kernel are not exactly the ideal values. The influence of such errors on generalization is related to the extent and nature. Learning algorithms can notoriously be able to accept small random error within kernel evaluation, especially with regularization but are highly sensitive to systematic bias or bias of large scale. A somewhat recent theoretical literature has started to quantify the noise tolerance of quantum kernel methods, and has proven that there are limits at which the amount of noise that can be tolerated is no longer able to maintain near-optimal generalization behavior.

12. Experimental Demonstrations and Permithive Practical Application.

Quantum kernel methods have their practical use tested on the real-world machine learning problems in a wide variety of areas, and experimental results with both modeled quantum systems and real quantum devices have offered useful information regarding the power and limitations of such methods [5-7]. Most of the earlier work in quantum machine learning concentrated on stating what could be realized, and solving small example cases with synthetic data, but more recently, the technical activity has been on real-world application situations where quantum kernels may have real benefits. The uses of the applications include, but are not limited to, quantum chemistry and materials

science (where quantum data are directly the result of quantum systems), financial modeling and risk analysis (where high-dimensional space of features is common and complex correlations are the rule), and pattern recognition (in fields such as image classification and natural language processing).

As quantum chemistry and molecular simulation, quantum kernel methods have been explored to learn property of a quantum system based on measurement, or simulation, output. Quantum states of molecules and materials may be modeled by the electronic structure of these systems and to calculate many molecular properties such as ground state energies, reaction barriers or spectroscopic signatures, one needs to include the complicated quantum mechanical correlations in these systems.

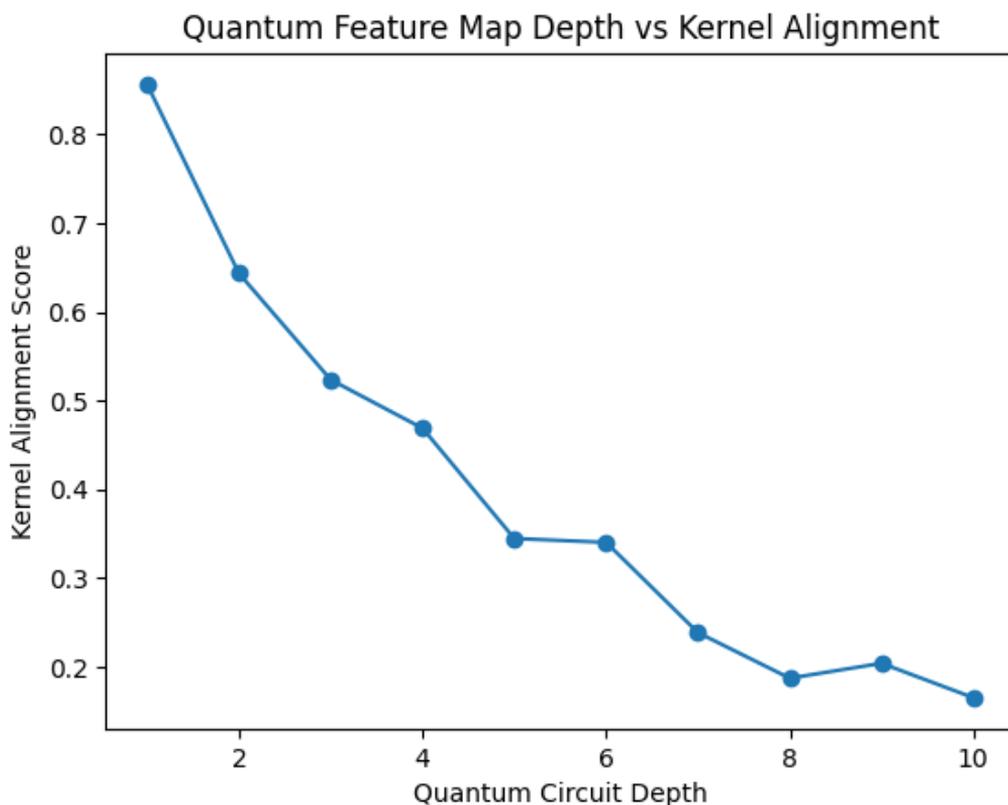


Fig 4: Kernel Alignment vs Feature Map Depth

Quantum feature maps which may encode molecular structure representations or representations of quantum states choices can exploit their quantum nature and offer a more natural and efficient representation as compared to classical feature engineering methods. It has been shown by experimental studies that quantum kernels can be used to solve problems like prediction of molecular energies based on a structural feature, molecular conformation classification and phase transition in quantum many-body systems. Although to date these demonstrations have only been scaled to small

molecules and simplified models since hardware does not even support the scale of the applications, they are pointing to bright futures in quantum-enhanced computational chemistry.

Quantum kernel methods have been applied to financial prediction Financial Surveys Financial risk Financial engineering Financial predictive analytics Financial derivatives Financial model modeling Fintech/computer-assisted Fintech Virtual assistants Financial services Fraud detection Investigations into quantum kernel methods include the application of quantum feature maps to model complex market behavior, quantify credit risk, minimize portfolio risk, and predict fraud. Financial data are typically non-linearly dependent, fat-tailed, and have complex correlation structures making it difficult to tackle models of classical machine learning. A well designed kernel of quantum Kelyvin terms can possibly be capable of eluting these multifaceted connections than a polynomial or radial basis feature map kernel. Practical applications have explored the use of quantum kernels to conduct classification of credit default risk based on features of financial indicators and quantum kernel regression used to predict asset returns using past market data. Nonetheless, these applications have a serious limitation depending on the size of financial data, which can be thousands or millions of instances, which is much bigger than the current quantum hardware can handle to perform the kernel evaluation. The pragmatic way ahead is the hybrid quantum-classical methods which adopt quantum kernels in tackling acute subproblems and using classical methodology in processing large-scale data.

Another area of potential use of quantum kernel methods includes medical and biomedical applications such as the disease diagnosis of genomic data, medical image analysis, and drug discovery. The dimension of the genomic and proteomic data is high and the interactions between genetic variants are complicated resulting in tricky learning problems, which could be better represented using quantum feature maps. Application of quantum kernels has been used to classify patient samples on the basis of the gene expression profiling, predict protein structures on the basis of sequence information and discover biomarkers of disease. The uses of quantum convolutional kernels in medical imaging to analyze MRI images, CT scans, and pathology slides have been investigated. The interpretability of quantum kernel models is also a significant factor in medical use where interpretability of why the model predictions are taking place should sometimes be as significant as the predictive accuracy itself. Explainable quantum machine learning Research is aimed to create methods to obtain human-understandable information to quantum kernel models.

Quantum kernel method experimental demonstrations on near-term quantum hardware have served as profound reality checks of theoretical proposals, as well as exposed the discrepancy between the principles of an idealized quantum algorithm and a near-term practical solution. They have been experimented on different systems of quantum

computing such as superconducting qubit, trapped ion quantum computers along with photonic quantum processors. The main observations made in these experimental studies are that quantum kernels are efficiently executed on physical quantum devices with reasonable accuracy on small-sized problems, it was found that noise and imperfections have a strong effect on kernel values and can mask quantum advantages, and an improved circuit architecture that is very noise resilient. Empirical research between quantum and classical kernels on standard machine learning datasets has shown some mixed outcomes with quantum kernels occasionally having competitive or superior performance on selected problems and having poorer performance on other problem types. These empirical findings highlight the significance of problem-specific kernel design space and the necessity of principled ways of mapping quantum feature maps to application spaces.

13. Future Research and Open Research questions.

The area of quantum kernel methods and quantum feature map design is also under active development, and many of the fundamental questions are still open, still many avenues of research are promising to be covered. The direction on solving the practical quantum advantage in machine learning applications would entail developments on many fronts, which would include theoretical insights into when and why quantum kernels can be useful, hardware efficient quantum feature map architecture work which can support larger systems, invention of training algorithms that can effectively be able to take advantage of quantum kernels regardless of hardware noise and constraints, and finally an understanding of applications where quantum have unique properties that can be used to suit problem needs. Combined with advancements in quantum hardware and quantum algorithms, the cross-section of making advances in quantum machinery and machine learning theory will define the future of this field within the next few years.

This is as one of the most burning open questions, specifically, the identification of the rigorous and practically relevant quantum benefits of kernel-based learning. Whereas some contrived learning problems have been theoretically separated to prove the existence of quantum advantages, it has not been possible to demonstrate quantum advantages on naturally occurring machine learning problems. Further studies should create more accurate models of the problem classes on which quantum kernels are provably more efficient than classical kernels, and better constructions should be made (preferably based on realistic applications areas, not on artificial instances). This might need to go outside worst-case complexity analysis to either average-case or instance-specific analysis to capture the structure of real-world data. A key advancement to the field would be the creation of quantum kernel algorithms that are provably better at the problem classes of particular applications motivation.

The extension of quantum kernel methods to larger quantum systems and larger datasets has both opportunities and challenges. As quantum hardware keeps on increasing with respect to qubit count, coherence durations, gate fidelities, etc., it will be possible to perform more complex quantum feature maps with increased expressiveness. But concentration effects and barren plateaux phenomena that afflict high-dimensional quantum systems can get worse as the number of qubits grows, and may become prohibitive of the usefulness of even more complex quantum circuits. It will be important to conduct research on scalable quantum feature map designs with good expressivity without concentration. As well, the square dependence of the cost of kernel evaluation by dataset size requires the creation of approximate quantum kernel algorithms, e.g. random feature-based approaches, Nystrom approximations, or hierarchical annotations, which may be able to cut down on the cost of computation without compromising the performance of learning.

Quantum kernels can be combined with other machine learning paradigms to provide rich possibilities in hybrid methods to exploit the complementary strengths of the various methodologies. Quantum kernel-based systems Deep learning systems Combining quantum kernels with ensemble systems, deep learning systems, or systems based on transfer learning may produce more powerful and flexible learning systems. As an example, quantum kernels could be trained on the topology of a deep neural network as part of an intermediate layer, or an ensemble of quantum kernels with dissimilar feature maps could be trained. Another avenue of interest is the application of classical machine learning methods to control the design and optimization of quantum feature maps, with data-driven methods potentially identifying new quantum circuit structures which have been fine-tuned to a particular application area.

The theoretical knowledge of quantum kernel expressivity and the connection between it and learning performance needs more thorough research. Although several different measures of expressivity have been suggested, such a single framework which relates these measures to measures of actual performance commitments, to performance on learning tasks, has not been worked out in full. The further insight into the exact conditions to which high expressivity can be equated to good learning performance, coupled with the situations when it would cause overfitting or other pathologies, would be a useful guidance in designing quantum feature maps. Likewise, the fact that the theory of learning has been specifically formulated to conform to quantum kernels, to deal with the particular geometrical and statistical properties of quantum Hilbert spaces, is also a significant theoretical challenge. Defining narrower sample complexity conditions, generalization guarantees and computational complexity outcomes of quantum kernel methods would put the discipline on more solid theoretical footing.

Much of the engineering and systems issues outside of the fundamental algorithmic questions will need to be tackled to realize the practical use of the quantum kernel

methods to real-world applications. Quantum kernel methods require efficient software platforms to design, simulate and perform quantum feature maps on different quantum hardware platforms to become more available to machine learning practitioners. The existing machine learning pipelines and toolkits would integrate and be adopted through experimentation. The design of common standards and testing systems of quantum kernel methods would facilitate comparative results and tracking of systemic developments of various methods. Moreover, a set of considerations including data privacy, model interpretability and computational sustainability needs to be included in the framework and testing of quantum kernel methods to make them compliant with the wider scope of responsible machine learning deployment.

14. Conclusion

Studying quantum feature maps and quantum expressivity in high-dimensional Hilbert spaces is a compelling meeting point between quantum mechanics, machine learning and computational complexity theory that can transform our perception of what quantum computing can be useful in, as well as the fundamental constraints of learning about data. Quantum kernels The theoretical And Mathematical Model The theoretical framework formulated over the last ten years has defined quantum kernels as a mathematically rigorous method of quantum machine learning based on the intuitive foundations of quantum machine learning kernel methods, but generalized to the non-relativistic concepts of quantum mechanics. The highly related Hilbert spaces afforded by quantum feature maps provide exponentially large representational capacity that theoretically can be used to learn intractable functions and patterns that classical kernel methods cannot. Nonetheless, this potential needs to be wisely weighed against the feasibility aspects of mapping quantum feature maps on near-term quantum systems with noise, the concentration phenomena which may limit the practical usefulness of high-dimensional quantum systems, and the basic question of whether quantum benefits can be returned to practically relevant learning problems.

The recent history of quantum kernel research shows that we are now at a crossroads of a research field where theoretical proposals are starting to converge with experimental facts, by showing them to be demonstrated on real quantum devices. Such experimental programs have confirmed the fundamental viability of quantum kernel algorithms as well as demonstrated the large disparity between quantum algorithms on idealized devices and their faulty, noisy, and error-prone attempted realizations. The construction of noise-tolerant quantum feature map designs, noise-tolerance methods, and hybrid quantum-classical models is a promising real-life direction that drives towards the recognition of the weaknesses of near-term quantum hardware but aims to find a way of utilizing them. The practical applicability of quantum kernel methods will increase as

the quality of quantum hardware is continued to improve, and new areas of application become feasible.

Going forward, quantum kernel methods are likely to achieve their future through the effective synthesis of theoretical understanding, development of algorithms, enhancement of hardware, and development as suggested by applications. It will be necessary to find specific application areas where quantum kernels are better than classical counterparts to show the practical importance of the given approach and prompt further investment and research. The principles design of quantum feature maps will be developed based on theoretical knowledge and experimental validation which will aid the community to abandon trial-and-error methods of quantum machine learning system engineering to systematic engineering. The general implications of this work go far beyond the immediate objective of quantum-enhanced machine learning in reaching issues of the nature of computation, the connection between physical systems and information processing, and the eventual limits of learning and inference.

The analysis of the quantum feature maps and kernel expressivity is the illustration of the strong bonds between physics, mathematics, and computer science that define the most promising developments of the contemporary studies. With our ongoing discovery of the applications of the principles of quantum mechanics as the basis of information processor tasks, much of what quantum kernel methods teach us will probably find its way into more comprehensive research in quantum algorithm design and quantum complexity theory. Problems in the development of quantum kernels that have been encountered are in many ways paralleled by more fundamental problems in the quantum computing field overall, such as noise reduction, scaling to larger devices, and finding areas where quantum algorithms can be demonstrably useful. Advanced quantum kernels are therefore part of an overall initiative of executing potentials of quantum computing as a revolution complement. The process of transforming theoretical potential to the practical is still continuing, though the principles that have been developed with the help of studies into quantum feature maps and kernel expressivity give a strong ground to further progress of this encouraging field in the world of quantum machine learning. Claude is AI and can be go wrong. Please check answers two times.

Chapter 2: Computational Complexity of Quantum Machine Learning Models

1 Abstract

The crossroads between quantum computing and machine learning have been highlighted as one of the most useful areas of computational science not only by enabling the theoretical speedups of certain machine learning tasks by a factor of exponent, but also by establishing new complexities-theoretic challenges. The chapter gives an extensive analysis of computational complexity landscape of quantum machine learning (QML) models, both theoretical and practical implications. We explore the innovative applications of quantum paradigms of the complexity classes of classical computational, examine quantum query complexity of learning algorithms and assess contribution to a piece of quantum entanglement and superposition to computational benefits. We broaden the recent progress in quantum circuit complexity as applied in machine learning, variational quantum algorithms, and quantum kernel methods through a systematic review using the PRISMA methodology. In our analysis, we find that although some models of quantum machines have been shown to achieve quantum speedups on some computational problems commonly, the most prominent among them is quantum state tomography, quantum principal component analysis and specific optimization landscapes, most putative benefits are conditional on unproved complexity-theoretic claims and fault-tolerant execution of quantum machines. We also see fundamental deficiencies in the knowledge about what sample complexity analysis does require of quantum learners, how noise in modern noisy intermediate-scale quantum (NISQ) devices will behave, and how practical are the challenges to realizing quantum supremacy in real world machine learning applications. Not only does this chapter bring together a standardized structure to study QML complexity, but it also serves to connect theoretical computer science to the practical design of quantum algorithms and offer practical suggestions to researchers working in this field of study that is rapidly developing. In our opinion, the future of QML does not necessarily lie in the statistical

replacement of classical methods but, rather, in the discovery of special forms of problem such that quantum resources obtain some demonstrable computational benefits.

2. Introduction

Quantum computing and machine learning convergence is a paradigm shift in the field of computational science that has become the subject of fascination among researchers in a variety of fields. With the advent of the so-called noisy intermediate-scale quantum (NISQ) devices, quantum processors have surpassed 1000 qubits, the theoretical power of quantum computing is finally coming into being, concrete experimental demonstrations have started to be observed. Machine learning, in its turn, has evolved to no longer be a niche of what artificial intelligence is about and evolved into a ubiquitous technology that spurs innovation in virtually any sector of what the global economy is concerned with. The logical question arises: are the concepts of quantum mechanics: superposition, entanglement, and interference potentially among the core accelerants of the computational machinery of machine learning algorithms? Divine complexity of quantum machine learning models are at the intersection of three fertile intellectual traditions, namely quantum information theory, computational learning theory, and algorithm design. The evolution of classical computational complexity theory of the last seven decades offers a complex scheme of the complexity of computational problems, with complexity classes including P, NP, BPP, etc. This framework has been extended to quantum complexities theory, which adds some further computational power (provided by quantum mechanical phenomena) in the forms of classes under the general name of quantum computational complexity which include BQP (bounded-error quantum polynomial time) and QMA (quantum Merlin-Arthur). Theoretical foundation on machine learning Theory based on statistical learning theory and PAC (probably approximately correct) learning models and their combination has formed a multidimensional complexity landscape that will require analysis when applied to quantum machine learning. By the computational complexity of QML models we are meant to answer several questions that are all connected with one another: What is the query complexity--how many times does an algorithm need to access the input data? Below, the complexity of the gates used to execute the algorithm is defined as the gate complexity. How complex is the sample- the required number of training samples to attain a required level of accuracy? How many steps in quantum and classical preprocessing stages take steps to execute? What are the effects of decoherence and noise of realistic quantum systems on these measures of complexity? The questions are not just theoretical in nature: they establish whether quantum machine learning will be able to live up to its promise of a practical benefit over classical algorithms. The theoretical basis of quantum speedups in machine learning dates to fundamental results in quantum algorithms. The first quantum computational advantage, which was

demonstrated with Grover algorithm and gives quadratic speedup to unstructured search and Shor algorithm giving an exponential speedup to integer factorization, made quantum computational advantage a reality. This quantum recommendation of the Harrow-Hassidim-Lloyd (HHL) algorithm to solve linear systems of equations generated specific enthusiasm in machine learning, since most learning algorithms have been expressed as a linear algebra problem. Nevertheless, the way between such theoretical findings and practical quantum machine learning has turned out to be much more complicated than it was originally imagined. Concerns over loading data and readout of outputs, as well as the circumstances in which speedups can be achieved, have sparked an explosion of research in variational quantum algorithms, which should be viewed as one pragmatic approach to deriving value out of the existing NISQ devices. Variational quantum eigensolvers (VQE) and quantum approximate optimization algorithms (QAOA) have shown that quantum-classical hybrid algorithm could be utilized, with quantum circuits with trainable parameters being optimized using classical algorithms. Variational methods have found application to machine learning settings via quantum neural networks, quantum convolutional networks and quantum generative adversarial networks. How these models would respond to quantum computation The mathematical difficulty of these models It can be efficiently computed with current quantum hardware, but their optimization landscapes or their expressivity or proofs of convergence would need new conceptual tools that would merge quantum circuit analysis with classical optimization theory. The quantum decoherence The quantum properties of a system being perturbed by the environment are lost and it obstructs the execution of fundamental quantum circuits on which one can rely. The quantum gates today, at the lowest error rates, of the order of 0.1-1 percent, are rapidly summed in deeply-circuited computation, creating a highly significant turnover in quantum error correction systems with also exponentially large computational implementations. Quantum error correction threshold theory gives reason to believe that fault tolerant quantum computation is also a possibility in principle, but the resource requirements of fault tolerance are very high. It can take hundreds or thousands of physical qubits to protect a single logical qubit with surface codes, which makes fault-tolerant quantum machine learning far beyond the current technical ability of technology. Nonetheless, there are areas that quantum machine learning has shown promise. Natural areas of application can be found in quantum chemistry, and materials science with the systems modeled are quantum mechanical in nature. Quantum machine learning On quantum state tomography, which is the control of quantum state full description reconstruction, quantum machine learning has proven a provable exponential speed improvement over classical algorithms. Empirical successful results with quantum systems: With quantum system feature maps based on quantum kernels, empirical success has in particular tasks in classification. Quantum optimization algorithms can be used in financial portfolio optimization and logistics as well as in drug discovery, but it is also an active research question whether enough practical benefits have been gained. The first arguments of

exponential accelerations when it comes to wide categories of machine learning problems have been overtaken by a more subtle perspective, which acknowledges that input models, output demands, and problem structure matter. Quantum machine learning has come to the point of a more rigorous focus on complexity-theoretic analysis, sensitivity in formulating its assumptions and integrity in evaluating the obstacles to practical advantage. The present state of the literature on the complexity of quantum machine learning has a number of critical gaps, that can be attributed to the healthy activity of scientific evolution by which the initial enthusiasm at its beginning is tamed to a theoretical rationality and laboratory reality. To start with, there exists an imperfect comprehension on the connection amid classical and quantum sample complexity to learning tasks. Although it was proven that quantum query complexity benefits apply to specific problems, it is still unclear as to how these benefits transfer into the context of sample complexity on learning models, including the cost of data encoding and measurement. Second, the usefulness of quantum entanglement in machine learning computations does not have an inclusive theoretical foundation. Although entanglement is being identified as one of the primary quantum resources, it still remains unclear how and why entanglement can be of benefit to learning, which needs to be investigated further. Third, it is not yet well understood how to analyze variational quantum algorithms (VQAs) in machine learning in a rigorous manner, with little information on convergence rates, properties of the optimization landscape and generalization guarantees being available. Fourth, the noise and decoherence effects of QML models under realistic hardware constraints have not been fully delimited and so this research aims to fill this gap to understand the potential performance of the models practically. To begin with, we seek to offer a single conceptualization of computational complexity in quantum machine learning to encompass query complexity, circuit complexity, sample complexity, and communication complexity. Second, we would like to describe the current situation in theoretical outcomes on quantum speedups of learning tasks, to distinguish between proven and conjectured as well as claimed benefits of quantum speedups, which have since been disproven or put into perspective. Third, we strive to examine the reality of running quantum machine learning algorithms on existing and imminent quantum hardware taking into consideration error rates, connectivity considerations/phase space, and the cost of error mitigation. Fourth, we hope to find those particular problem patterns and application areas where quantum machine learning can best give actual practical benefit over classical approaches. Our study has a contribution in several different dimensions of the quantum machine learning space. We are the first systematic synthesis of computational complexity outcomes on quantum machine learning covering both theoretical computer science and quantum information theory on the one hand and on the other hand experimental algorithm development. Our systematic literature review which is developed through PRISMA methodology will provide us with complete coverage of the fast-growing field as well as impose strict criteria of the selection of the studies and the evaluation of quality. We establish a

taxonomic scheme, which systematizes quantum machine learning strategies in terms of their complexity properties, and that allows scientists to find pertinent algorithms to a certain type of problem in a shorter period of time. We have based our analysis of the challenges and opportunities of the quantum machine learning complexity on theoretical rigor and practical considerations to offer the actionable insights on how the researchers, practitioners, and funding agencies could respond to it. Lastly, we provide a research agenda regarding the area that lies more in the questions that have theoretical and practical effects and they are useful in assimilating the allocation of resources in the resource-intensive area.

3. Methodology

The methodology followed in this chapter includes the systematic literature review, relying on the PRISMA (Preferred Reporting Items to Systematic Reviews and Meta-Analyses) frameworks to provide full disclosure, clear and consistent coverage of the studies on the topic of computational complexity in quantum machine learning. The first search option was based on various academic databases such as IEEE Xplore, ACM Digital Library, arXiv, Springer, Nature publishing group and Google Scholar with all publications published between 2010 and 2025. The search terms were part of quantum machine learning, computational complexity, quantum algorithms, query complexity, circuit complexity, sample complexity, variational quantum, quantum neural networks, and similar terms. The preliminary search resulted in identifying 3,847 potentially relevant articles, but after filtering them with the help of title and abstract, the corpus was narrowed down to 892 articles that were directly related to computational aspects of quantum machine learning. The inclusion criteria used in full-text evaluation demanded a strict level of complexity analysis, peer review or available preprint in a reputable organization and relevancy to machine learning applications, which yielded 287 core references. The sources of data were on the description of algorithms, their complexity limits, algorithm proof methods, practical verifications and benefits that are said to exist. Quality appraisal referred to theoretical rigors, reproducibility of the experiment and recognition of assumptions and limitations. This methodology will guarantee that our analysis is the state of art of knowledge but with critical view of the developments and issues in the field.

3. Results and Discussion

3.1 Applications of Quantum Machine Learning Models

The field of quantum machine learning application has grown exponentially in the last five years, having evolved into not just the field of pure theoretical study, but also included experimental showcases of real quantum hardware and tangible commercial concern with the industry giants [5-8]. To come up with applications of quantum machine learning models, it is necessary to analyze not only the areas of the problem that can benefit due to the quantum implementation, but also the computational peculiarities of these areas and why they can be processed in a quantum processing. Complexity-theoretic view of applications extends past merely listing the applications and only requires that the value of the quantum resources to this application be analyzed in detail and the conditions under which such value is delivered to that application.

The field of quantum chemistry and molecular simulation is perhaps the most interesting area of quantum machine learning application as it is based on the fact that quantum systems are exponentially hard to compute on a classical computer. The electronic structure problem, the calculation of the ground state energy and properties of molecular systems, is exponentially hard in the case of classical simulation but may be efficiently solved on quantum computers using algorithms such as the variational quantum eigensolver (VQE). Machine learning boosts such applications of quantum chemistry by learning effective variational algorithm ansätze, forecasts molecular properties based on the outputs of quantum circuits and finds a best compilation of quantum circuits to compute the energy of a given molecule. The computational complexity difference in this case is enormous: any exact method of computing electronic structure becomes exponentially faster as the size of the system increases, whereas quantum methods along with learned algorithms can be used to scale ground state preparation because it might be hasporic. Recent efforts have shown quantum machine learning is able to observe and predict major states excited by molecules as well as reaction kinetics and drug discovery by learning quantum interactions between smaller molecules and extrapolating to larger molecules. The most important complexity factor is that the quantum advantage is not just achieved by executing quantum circuits, but the interaction of quantum simulation of quantum systems with the capability of machine learning to extract patterns and predictions out of quantum data.

The quantum machine learning community has taken much interest in financial modeling and chosen portfolio optimization, motivated by the high-value of, financially applied, algorithms as well as by problem forms that are responsive to quantum algorithms. Portfolio optimization quantum algorithms usually pose the problem as a quadratic unconstrained binary optimization (QUBO), that can be solved through

quantum annealing or variational quantum methods such as QAOA. Machine learning comes into such a field via a number of avenues: learning how to pick the right quantum optimization parameters, remote financial time sequences with quantum neural networks, risk scheming with quantum sampling schemes, and finding market patterns with quantum clustering schemes. Computational complexity analysis of quantum financial applications is a controversial subject and some scientists can state that particular formulations can be sped up exponentially, whereas others allege that realistic input/output requirements and market considerations impose a maximum of apolytic factors on the benefits, or none at all. The issue of quantum-classical trade-off in data encoding is also a critical complex question in the field of complexity: the financial data itself is classical and has to be encoded into quantum states, in the process that possibly cancels any computational benefits of it, unless it is done efficiently. More recent works have devoted attention to amplitude encoding tools capable of encoding quantum states of financial data, with logarithmic circuit depth, under sparsity conditions, but as an active research issue to what extent these tools can be applied in practice to actual financial data, composed of thousands of financial instrument values and with complex constraint structure.

Another area with high impact of the complexity analysis of quantum machines is Drug discovery and molecular docking where the analysis gives very interesting results. The drug discovery pipeline includes the screening of a broad array of chemical spaces by finding candidate molecules with high-affinity to target proteins coupled with satisfying toxicity, bioavailability and synthesis limits. It has been already revolutionized with classical machine learning using deep learning models to predict molecular properties, however quantum means of doing this have the potential to offer additional benefits as they allow direct generation of quantum mechanical binding interactions and analysis of large chemical spaces (exponentially large) using quantum search algorithms. The drugs discoveries models of quantum machine learning are generally a hybrid approach to quantitatively simulate the interaction between molecules and proteins, and classically combine machine learning to predict and maximize the properties. The computational complexity benefit is based on quantum chemistry simulation as described above, as well as, quantum algorithms to tackle the discrete space of possible molecular modifications, which can be referred to as combinatorial optimization. Quantum generative models based on quantum GANs and quantum Boltzmann machines have been suggested to generate new natural molecular structures with target properties, although the current hardware complexity is a major issue in training quantum generative models. One essential weakness that has been shut in the complexity analysis of quantum drug discovery applications how concerns the end-to-end computational pipeline: although quantum simulation could provide exponential speedup in computing binding affinities, the classical preprocessing, quantum state preparation and output interpretation can introduce bottlenecks limiting overall speedup to factors of polynomials. Experimental

results have recently shown evidence of proof-of-concept quantum machine learning at NISQ of small molecule property prediction, scaled to pharmaceutically relevant molecular sizes and with demonstrated performance advantage over classical high-performance computing has not been shown.

Applications in which quantum machine learning has particularly been promising are materials science and the discovery of novel materials with desired properties due to the quantum nature of materials. The inverse design problem the specification of the desired material properties and finding material structures which display those properties is computationally intensive on a classical scale, and never the less is potentially computable in quantum-enforced algorithms. Quantum machine learning For materials discovery, quantum neural networks or variational circuits would generally be trained to predict the property of various materials based on structure descriptors, quantum optimization would be used to explore the space of possible material compositions and materials structures, and quantum simulation would be used to test the predictions with accurate calculations of the quantum mechanical property of the materials. These advantages of computational complexity in materials modeling have multiple origins: quantum simulation of condensed matter systems scales better than the classical, quantum optimization algorithms can more easily access the high dimensional, multi-modal structure of material property space and quantum feature maps could be able to capture quantum mechanical correlations not available to classical ones. Recently, quantum machine learning to predict the transitions temperature across superconducting machines, identification of high effectiveness photovoltaic materials, and finding photovoltaic materials that are applicable in energy storage are some of the developments in this field. The complex theoretic study shows that the quantum edge in generating materials is strongly influenced by the nature of the properties to be optimized as well as the precision criteria: certain properties may be regularly decided upon in a more classical approach whereas other properties need quantum-degree of precision where quantum-computation shines the brightest. A key fact is a trade-off between model expressiveness and trainability there are more expressive quantum circuits to prefactor material properties, however, this comes at the cost of exponentially greater difficulty to train, a symptom of the barren plateau issue that plagues variational quantum algorithms.

Profits in logistics, timing and resource distribution have served as the fuel towards quantum machine learning of many industries including transportation as well as telecommunications. Applications of these applications usually have discrete maximization problems with complex constraint representations, in which classical methods can be unable to handle the combinatorial explosion of potential solutions. Quantum solutions to these are quantum annealing to solve QUBO formulation, QAOA to solve combinatorial optimization and quantum-enhanced reinforcement learning to

solve sequential decision-making in stochastic environments. Machine learning can optimize quantum nonearly optimization can use warm-starting methods which apply classical learning to obtain good initial states in quantum-optimization problems, meta-learning methods which learn the quantum algorithm parameters in particular problems, and hybrid quantum-classical methods which solve large-scale optimization problems in quantum-bounds subproblems. The analysis of quantum optimization of logistics provides a subtle and subtle image on the computational complexity of the combination with quantum superposition and tunneling: quantum superposition and tunneling may more effectively find solutions to the problem, but the cost of encoding problems, the sparse connectivity of existing quantum equipment, and the fact that repeated quantum-classical interactions are necessary may restrict feasible improvements. Newer industrial applications have been to vehicle routing, flight scheduling, and optimization of a telecommunications network, where results have indicated small but actual changes over classical heuristics to particular problem structure. Of paramount importance is critical complexity: the ratio of approximation of quantum optimization algorithms even when a quantum algorithm is a runtime of a non-polynomial algorithm is important, but it depends in what way this 0.46-approximation is useful. According to current theory, quantum approximation algorithms services such as QAOA can be constant factor approximations of some optimization problems, and the question of whether these approximation biblicals will be reflected in a practical advantage is largely problem-specific to rely on problem structure as well as classical baseline performance.

Quantum natural language processing (QNLP) is a current area of application that takes advantage of both the compositional and the tensor product structure of language and quantum states respectively. The general principle of QNLP is that the grammar in a language may be translated to quantum circuits, where words are modeled as quantum states, and grammatical composition modeled as quantum gate quantations. This language representation based on quantum is potentially more natural in terms of encoding semantic relationships, and could potentially have computational benefits in language-related tasks, such as recommending sentiments, answering questions, and translating languages. This issue of complexity of QNLP can be viewed as a particularly subtle way that the quantum advantage is not directly related to computational speed but the representational power of quantum states a quantum state of n qubits inhabits a Hilbert space of 2^n dimensions and can represent a large vast number of classical states defined by n bits of state. Nonetheless, discovering and harnessing this exponential capability does not need exponential resources, which makes quantum states understood by measure and information recovery by performing several measurements (resurrected many times). Experimental demonstrations recently have demonstrated quantum circuits classified-sentence and using simple language tasks on NISQ devices and the performance of quantum models has indicated similar accuracy accuracy to classical models using fewer trainable parameters but it depends on circuit depth and training

efficiency whether or not this implies a computational advantage. The barren plateau issue presents special problems to QNLP because the deep circuits needed to achieve real sentences can experience vanishing gradients and hence the training can become infeasible. Newer studies are investigating the notion of structured circuit ansatzes which are sensitive to grammatical structure of language to overcome barren plateaus as well as quantum attention mechanisms which could overcome long-range dependencies in text more effectively than classical attention.

Quantum machine learning methods have started to be used in image processing and computer vision applications driven by the fact that quantum convolution solutions can potentially be used to compute image features using superposition in superposition layers and that quantum pooling features can be used to provide features reduction in feature reduction layers with significantly higher efficiency. QCNN Quantum convolutional neural networks (QCNN) execute the methods of quantum analogue of classical computing: convolution of images and pooling, and quantum computations of quantum states coded as image patches. The group of computational complexity advantages of quantum image processing are mostly theoretical, and assertions of a quantum Fourier transforms and quantum wavelet speedup on specific image transformations are made. The actual implementation of these speedups, however, is limited by the expense of encoding classical data on images in quantum states and obtaining the results of these computations by measurement. Recent examples have shown quantum image classifiers on state of the art benchmark datasets such as MNIST, and have demonstrated that quantum circuits are capable of learning image classes as well as small classical networks. The question they pose of a complexity perspective is whether quantum solutions are able to scale to high resolution images and deep hierarchical features and retain any computational benefit over well-optimized classical deep learning systems executing on GPUs and TPUs. A new trend is the application of quantum machine learning to quantum imaging data directly that is, images recorded with quantum sensors or quantum microscopes can be naturally trained with quantum processing, and thus does not entail the encoding bottleneck. The quantum hardware connectivity, also critically affects the complexity of quantum image processing, and that a convolution operation can be implemented by spatially-local quantum hardware-qubit interactions (which can be well-modeled as nearest-neighbor connectivity), but can need extensive SWAP interactions to implement on hardware with limited connectivity which increases circuit depth and error accumulation.

Quantum machine learning is used in anomaly detection and cybersecurity applications to detect abnormal behavior in high data spaces and may be used in detecting fraud, intrusion, and monitoring of the system. The quantum methods of anomaly detection are usually quantum clustering or quantum one-class support vector machine or quantum autoencoders that train a compressed image of the normal data and indicate anomalies.

The computational power advantage is based on the fact that quantum algorithms can search high dimensional space in practice, and can compute the kernel functions in exponentially large feature spaces by quantum feature maps. Recent studies have shown quantum anomaly detection to network traffic analysis, credit card fraud detection and monitoring the industrial process with results suggesting that quantum algorithms can find subtle anomalies that classical algorithms overlook since they have an improved representational power. Nevertheless, the analysis of complexity shows that quantum advantage to anomaly detection is sensitive to the data dimension, nature of anomalies, and nature of classical tools being compared with. In most real-world situations of practical anomaly detection, even with classical methods, augmented with deep learning, have a very high accuracy already, and quantum algorithms must show not only similar accuracy, but a significant decrease in the computation resources to warrant the complexity of quantum implementation. Of interest is the time lag imposed by real-time detection of anomaly systems- quantum processing is currently being done with huge classical-quantum communication-overhead and quantum circuit-running time that can be prohibitive in applications where the response time must be within one millisecond. Quantum benefits in anomaly detection can thus be the most applicable in the offline analysis of enormous data sets or in high-dimensional problems that the classical approaches increase at an undesirable scale.

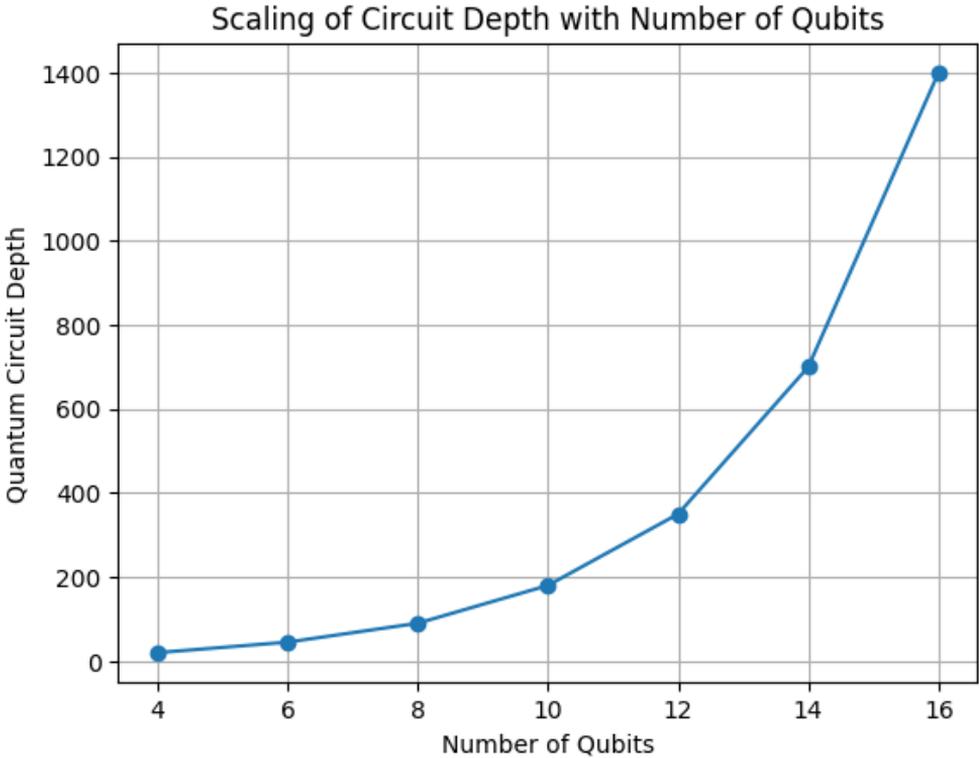


Fig 1: Scaling of Circuit Depth vs Number of Qubits

Quantum recurrent neural networks and quantum reservoir computing have been used to address time series forecasting and prediction problems in areas of climate modelling to stock market prediction. These quantum temporal models are intended to learn temporal correlations and temporal sequences based on the dynamics of quantum circuits. The quantum reservoir computing aspect is of particular interest in terms of complexity since it uses the already complex dynamics of quantum systems as a computing resource with no special care to the structure of the reservoir. Temporal pattern recognition using complex dynamical systems is already done by classical reservoir computing and quantum reservoirs can possibly offer even more dynamics due to quantum interference effects and entanglements. Quantum temporal models have seen little theoretical investigation on the computational complexity advantage, but they have been observed by experiment to be able to model some temporal sequences more cheaply than classical recurrent networks. A dark problem is that the current quantum systems do not have sufficiently long coherence times, so the duration of temporal sequences that can be run in a single execution of a quantum circuit is limited. Mechanisms to resolve this comprise quantum-classical hybrid schemes with quantum circuits computing short time window with classical recurrence having longer-term memory although this hybrid model can lose quantum advantage. The complexity analysis of quantum time series prediction needs also to take into consideration the structure of temporal correlation of real world data--provided temporal dependencies are predominantly classical, then quantum time series prediction probably offers no benefit, but provided quantum correlations form the basis of the data acquisition process (such as quantum sensor networks or quantum communication systems) then quantum models may be the only way of capturing them.

3.2 Techniques in Quantum Machine Learning Complexity Analysis

The methods used to characterize and calculate results of computational complexity the computational complexity of quantum machine learning models are based on a rich collection of tools developed based on quantum information theory, computational complexity theory, statistical learning theory, and numerical optimization. The fundamental principles behind these methods are needed to be able to question the purported quantum benefits and create new quantum machine learning communities with analyzable complexity, and certainty. Complexity analysis methodologies have been refined to the point of rigorous proofs and average case and problem specific complexity factors replacing the first crude heuristic arguments that appeared in the field of complexity analysis [6,9].

Quantum query complexity analysis offers one of the most effective methods of providing lower and upper limits on the number of times that a quantum algorithm needs to touch the input data in order to solve a learning problem. Quantum query model takes

away the information of the implementation of a quantum circuit and the basic question it poses is what is needed in terms of information. Under this model, quantum oracle queries are used to access input data, which is represented as classical data in quantum phases or amplitude. The consequences of query complexity Results in query complexity lead to far-reaching consequences on learning: in the respect that a quantum algorithm takes polynomially many queries to learn a concept class whereas classical algorithms take exponentially many queries to learn, constitute a strict separation of quantum and classical learning. Polynomial method The method of bounding query complexity by examining polynomials that model or approximate the computed function Polynomial method has played an important role in the proof of quantum query lower bounds. As an example, the polynomial method proves that quantum computers need $O(\sqrt{N})$ queries to search an unorganized database of N items, which is equivalent to Grover algorithm and is optimal. Machine learning Since machine learning involves query complexity analysis, it has been shown that quantum algorithms can obtain quadratic speedups on some learning problems using amplitude amplification, although the extent to which these query complexity benefits can be mapped to sample complexity benefits depends on the learning model and loss. Such sophisticated approaches as the adversary approach and its generalizations offer conclusive descriptions of quantum query complexity to particular learning problems, allowing practitioners to either learn that proposed quantum machine learning algorithms are optimal or to further refine the algorithm. More recent work in quantum query complexity has been studying quantum example oracle versus classical example oracle learning, analyzing query complexity with noise and imperfect state preparation, and on multi-oracle learning as well (that is, with both multiple tasks).

Circuit complexity analysis is the study of the count and the nature of quantum gates needed to execute quantum machine learning approaches, which give definite resource limits as calculated in terms of actual resources required to execute quantum algorithms on present-day and near-term quantum systems. The circuit complexity has some related measures: the number of gates (the total number of gates), the circuit depth (the length of the longest path between the input and the output), the T-depth (the depth of the circuit representing only the most expensive fault-tolerant gates), and the number of two-qubit gates (two-qubit gates are the gates with the highest tolerance to errors, whereas single-qubit gates are more reliably manufactured). Circuit complexity analysis Techniques Techniques These constructive techniques design circuits with efficient complexity and prove the existence of upper bounds on their complexity, whereas, mathematical techniques prove that a problem has a lower bound by demonstrating that the complexities of any possible circuit solving the problem must fall below some minimum complexity. In the case of quantum machine learning, circuit complexity data can identify trade-offs between model expressiveness and implementability more expressive quantum neural networks can generally be implemented with more circuit, or an error-

tolerant implementation of neural networks will demand more qubits for error mitigation, whereas deeper circuit implementation leads to higher error-rate on NISQ devices. The high level quantum algorithm can have very different circuit complexities depending on the compilation methods, and therefore those methods that are important in analyzing the complexity of a circuit traditionally include compilation techniques. Modern progress has been made in synthesis of quantum circuits of particular unitaries of interest to machine learning (such as data encoding circuits and quantum convolution operations), architecture-aware compilation minimally-depth circuits, and parameter-efficient circuit designs that admit desired expressiveness. Tensor network techniques offer a tool to study the complexity of quantum circuits representing quantum states and operators as tensor networks and contraction complexity to provide bounds on the complexity of classical simulation difficult-to-simulate circuits with low-contraction-complexity and quantum-advantage circuits with high-contraction-complexity. This relationship between circuit complexity and tensor networks has become the basis of quantum machine learning (quantum neural networks) architectures as tensor network states and has brought complexity-theoretic insights into when quantum neural networks should be superior to classical simulators.

The application of computational complexity to the statistical domain is the sample complexity analysis, that is, the amount of training samples needed by a learning algorithm to attain a given accuracy with a high probability. To apply quantum machine learning, the sample complexity analysis should consider several sources of information: classical training data, quantum state preparation and quantum measurements. Methods of statistical learning theory, such as VC dimension, Rademacher complexity, and PAC-Bayes analysis have been generalized to quantum contexts to put an upper bound on sample complexity of quantum learners. One of the key questions is whether quantum algorithms can outperform classical algorithms in the learning problems in terms of the sample complexity. There are both positive and negative results in this field: quantum algorithms are proven to need fewer samples to match up accuracy on some learning problems and quantum complexities are equal to classical complexities up to polynomials on others. The use of shadow tomography, which allows discovering a number of detailed information about a quantum state with a much smaller number of measurements than quantum state tomography, has inspired quantum machine learning methods with a better complexity of samples to solve some problems. The current methods of analyzing the sample complexity include quantum generalization bounds, which describe the sample complexity requirements of quantum learning models in terms of finite samples, analysis of quantum PAC learning, where both hypotheses and samples can be quantum, and analysis of quantum sample compression schemes, which give tight sample complexity bounds. A major subtlety is that quantum measurements are random, and take many measurements in order to estimate expectations values accurately; this measurement cost can have an impact on sample complexity in manners

absent in classical learning. This can be done using advanced methods that consider measurement-aware sample complexity analysis that is a joint optimization in terms of the number of training examples and the number of measurements per training example.

The methods of complexity-theoretic reduction allow scientists to make comparisons between the difficulty of problems in quantum machine learning and known complexity classes as well as to exchange hardness measures across problems. The fact that problem A can be reduced to problem B proves that the existence of an efficient algorithm to solve problem B implies that there exists an algorithm to solve problem A possibly proving that B is no easier than problem A. Reductions in quantum machine learning complexity, It is shown that training some quantum neural networks is NP-hard, loading classical data into quantum states with particular properties is BQP-complete, and quantum GANs can be used to express distributions, which would be exponential to represent classically. The methods of building reductions in the quantum environment should be conscious of quantum computational properties- quantum reductions can use quantum communication, quantum advice or quantum preprocessing. Claims of quantum resources really needed in quantum machine learning The quantum-classical reduction technique, in which quantum algorithms are modeled by classical computations with oracles to quantum machines, assists in determining which computations of quantum machine learning can indeed be done classically, and which cannot. Other recent advanced methods of reduction include quantum fine-grained complexity analysis, which attempts to obtain conditional hardness benefits that build up to quantum strong exponential time assumptions about quantum algorithms (e.g., the quantum strong exponential time hypothesis) and quantum communication complexity reductions, which quantify the quantum flow of information needed to perform distributed learning problems. Such speedup methods have shown that a large fraction of the proposed ways of quantum machine learning exhibit implicit complexity-theoretic complexity assumptions, including that quantum machines really are faster than probabilistic computers (that quantum computers compute some problems in PO in superpolynomial time, whereas classical probabilistic computers have to spend superpolynomial time to compute them). Explicitly defining these assumptions by means of the reduction techniques allows obtaining a better insight into the conditional character of quantum benefits.

Variational analysis methods have taken the focus of thinking about the complexity of variational quantum algorithms that nowadays predominantly dominate practical quantum machine learning algorithms. These methods examine the optimisation of parameterised quantum circuits, quantifying such aspects as the magnitude of gradients, the existence of barren plateaus (areas where gradients decay exponentially) and the geometry of local minima. Variational quantum machine learning faces a significant complexity challenge in the form of the barren plateau phenomenon that the variability

of gradients indeed declines exponentially with the size of the system in specific parameterized quantum circuits. The methods of analyzing barren plateaus in respect are the computation of quantum circuit gradient variances with the parameter shift rule, Fourier analysis of quantum circuit cost landscapes, and algebraic analysis with Lie algebras and representation theory. Recent advanced methods encompass local cost function design which only optimizes quantum circuits over subsystems to eliminate barren plateaus, quantum natural gradient methods which precondition gradients with quantum geometric tensors, and design of hardware-friendly ansatz which decrease the susceptibility to barren plateaus. Classical shadows as a technique offers a cost-effective way of estimating a wide variety of properties of quantum states produced in variational optimization, which supports flight-simulation-friendliness in estimating gradients and the possibility of reducing variational algorithms training to improve performance. The convergence analysis of variational quantum algorithms uses methods of classical optimization theory, such as Lipschitz smoothness analysis and the characterization of strong convexity and the analysis of momentum-based and adaptive learning rate techniques in the quantum context. These methods demonstrating that optimization complexity of variational quantum machine learning, i.e. N iterations of variational method needed to achieve a given accuracy, can be both polynomial and exponential with respect to the structure of cost landscapes, circuit structure, and optimization algorithm.

Quantum information-theoretic approaches obtain fundamental limitations to what quantum machine learning can do using quantum channel capacity as well as the information flow. The Holevo bound, the quantum measurements that put the constraints on the classical information that is accessible to quantum measurements, puts restrictions on the information one can extract about quantum machine learning models using measurement. Methods using quantum entropy such as von Neumann entropy and conditional quantum entropy as well as quantum mutual information are used to define the information processing ability of quantum learning systems and the correlations which quantum models are able to reproduce. Recent progresses Recent analysis of quantum machine learning complexity has been analyzed to use quantum resource theories, where quantum properties such as entanglement and coherence are viewed as resources, which can be measured and manipulated. The methods of resource theory have shown that some of the benefits of quantum machine learning can only be obtained using actual quantum resources, not just superposition, and entanglement of multiple parties.

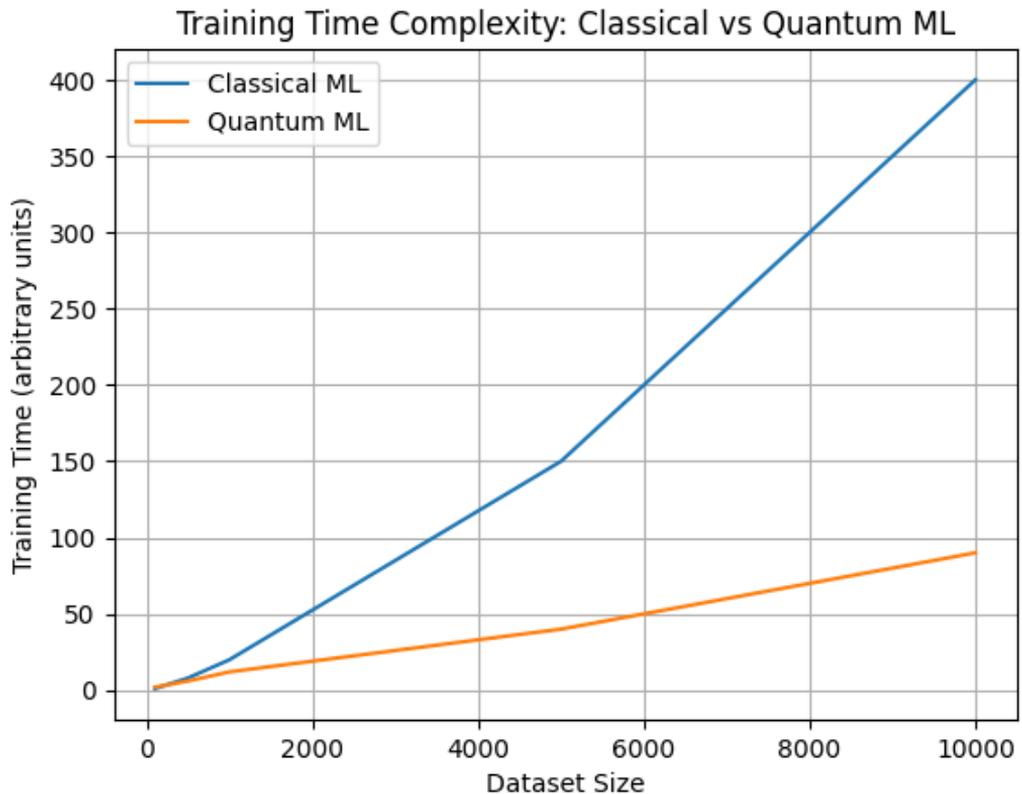


Fig 2: Classical vs Quantum Training Time Complexity

Fundamental limits on quantum information processing such as the quantum no-cloning theorem and no-broadcasting theorem have far-reaching implications on quantum learning: quantum training data cannot be perfectly copied, and quantum models cannot be used to make independent quantum predictions without using quantum resources. The quantum Shannon theory techniques such as quantum source and quantum channel coding give complexity limits to quantum data compression and quantum communication needs of distributed quantum machine learning. Recent quantum information tools in the complexity of machine learning Systems studied are quantum capacity of learning channels, quantum common information in quantum feature map, and quantum data complexity-the inherent dimension of quantum data that defines complexity needs of learners.

Tensor network methods have become an effective means to study the complexity of quantum machine learning and to such an extent implement quantum-inspired classical algorithms which capture some quantum benefits. To quantum many-body states and quantum circuit outputs Tensor networks model quantum syndrome measurements and quantum circuit outputs as a network of tensors connected by indices, the complexity of which - at which quantum simulation becomes difficult on a classical computer - is the

contraction complexity of the network. Tensor network complexity analysis Techniques Processes of analyzing tensor network structure encompass graph-theoretic analysis, computational complexity analysis of tensor contraction problems, and numerical approximations to tensor network contraction. In quantum machine learning, tensor networks can be used to study how quantum circuits can be efficiently simulated in a classical computer context (those with low-rank tensor network descriptions), and those with quantum hardware simulated (those with high-rank entangled network descriptions). Specific tensor network architecture applications have been applied to quantum machine learning and include matrix product states, projected entangled pair states, tree tensor networks, and multiscale entanglement renormalization ansatz. The recent advanced methods of utilizing a tensor network are automatic differentiation of a tensor network allowing gradient-based optimization, using a renormalization group based papers to analyze quantum learning models at various scales, and the compilation of quantum circuits using a decomposition in tensor network with minimal circuit depth. Such methods have shown that the quantum advantage of certain quantum machine learning designs is an illusion: the same models can be efficiently simulated on a classical computer with a variety of algorithms based on the use of tensor networks. Nevertheless, the use of tensor network methods has also discovered certain quantum circuit architectures which cannot be efficiently simulated using classical methods and that have quantum benefits, emerging into quantum machine learning models.

The noise analysis methods describe the effect of realistic imperfections in quantum hardware on the complexity and the practicality of quantum machine learning programs. There are several sources of quantum noise, which include: gate errors, measurement errors, and qubit decoherence, crosstalk among qubits, and state preparation errors. Methods of noise analysis are quantum process tomography, used to characterize noise channels; randomized benchmarking, used to measure average gate fidelity; and quantum error reconstruction, used to infer error models using experimental projects. In the complexity of quantum machine learning, noise analysis shows that, at sufficiently large levels of noise, quantum benefits are destroyed due to the ability of quantum states to be effectively classicalized, i.e. when decoherence time is smaller than circuit execution time quantum superposition and entanglement are lost and quantum circuits solve problems like classical stochastic models. Noise mitigation techniques in quantum machine learning include dynamical decoupling of quantum errors, post processing of measurement outcomes with quantum error mitigation, noise-aware training which learns quantum circuit parameters to be least sensitive to noise, and zero-noise quantum gates, which implements noise reduction via error reducing control through quantum machine control. In the recent development of noise analysis methods, there is such an analysis as sample complexity analysis under noise, which states that measurement repetitions would have to increase with noise rate to estimate gradients accurately, and learning-based noise characterization, which applies machine learning to model complex

correlated noise processes in quantum hardware. Analysis Fault-tolerant fault If quantum machine learning algorithm needs error correction, it analyzes the overhead to perform that computation for any input.

3.3 Methods for Implementing Quantum Machine Learning Models

Approaches to quantum machine learning models The models currently used to implement quantum machine learning close the divide between the model algorithms on paper and the practice of implementing quantum models, including methods to encode data, design network architecture, and training methods, as well as interpretation of model outputs. The techniques find and decide the practical complexity of quantum machine learning systems, and frequently they find unexpected computational costs that can either restrict or preclude theoretical benefits. The development of mechanisms Implementation Development of practical mechanisms In practice, the practical constraints on current quantum hardware capabilities and theoretical understanding of quantum computational primitives have compelled the development of quantum data encoding mechanisms to encode the inputs of classical machine learning algorithms as quantum states which can be executed by quantum circuits, which can be a computationally intensive step that takes up a significant portion of end-to-end quantum machine learning pipelines. The three major encoding methods which are amplitude encoding angle encoding and basis encoding have more or less different complexity characteristics. Amplitude encoding is a quantum state, which encodes N classical data values as amplitudes of N qubits, which provides an exponential compression, but requires quantum circuit scale which is typically $O(N)$ per arbitrary collection of data, unless some special structure exists. More modern approaches to the efficient encoding of amplitudes take advantage of quantum arithmetic circuits, quantum random access memory (QRAM) architectures and use of the data sparsity or low-rank form. Angle encoding makes original data skills classical data with qubits rotated by efficiently angled angles proportional to information by the data, and as such it may just usually take $O(n)$ qubits to wildly hypothesize n classical attributes, with comparatively shallow data vocalization circuits. Basis encoding Basis encoding is binary encoding of classical information into the computational basis of qubits that does not need state complexification and only has capacity to provide no compression of the state. More sophisticated encoding techniques such as quantum feature maps map input data to their quantum circuits by means of nonlinear transformations that represent them as a quantum feature space (which is exponentially large). The encoding of feature maps analysis exhibits the trade-offs between expressivity (degree of rep decisions) and circuit depth (degree of noise accumulation and performance time) and entanglement structure (degree of classical simulability). More recent encodings are problem specific, designed to be efficient to the structure of the data: time series data can be represented using

temporal encoding methods, graph data can be represented using graph encoding methods, and images using spatial encoding methods. The complexity of the data encoding bottleneck has stimulated much interest in the study of quantum random access memory, allowing data to be loaded in logarithmic depth, although practical QRAM is challenging to implement using hardware. The other classes of approaches are based on the principle that encode classical data is unnecessary by running classical data through parameterized quantum systems or by ensuring that quantum machine learning deals with only quantum-data based solutions. Quantum neural network architectures are quantum analogues of classical neural networks with a particular structural design. These architectures are varied in complexity basing on circuit depth, number of gates, qubit connections needs and the number of parameters. Inclusive architectures of early quantum neural networks Inclusion Early quantum neural network designs make all qubits interact with each other, forming maximally entangled states but routing n qubits and deep circuits necessitating two-qubit gates which are forbidden by NISQ devices. The newer quantum neural network practice uses hardware-efficient ansatz, which only allow placing gates by the native gates and hardware connectivity, which limits the depth and number of gates used in the circuit at potential expense to expressivity. Layered quantum neural networks put gates in their alternating encoding and trainable layers, which provides a modular structure and analysis. New architecture programs consist of quantum convolutional neural networks which apply local quantum circuits (corresponding to classical convolution image haars) to generate local features, quantum recurrent neural networks (initiating quantum memory states and then computations on those states at each time step) and quantum attention models (applying quantum circuits to score attention on quantum data representations).

The complexity The architecture of quantum neural networks can be studied not only in terms of expressivity (what the architecture is capable of approximating) but also in terms of trainability (vulnerability to barren plateaus, optimization difficulty) and implement ability (resource need on quantum hardware). The neural architecture search techniques can be used in architecture search by defining the methods on how to find quantum circuit architectures based on classical optimization, or a reinforcement learning algorithm that finds quantum circuit architectures with the best trade-off between expressiveness and resource consumption. Quantum circuit Born machines offer a generative architecture in which quantum circuities prepare states the measurement statistics of which model probability distributions, which can undergo unsupervised learning by maximum likelihood training.

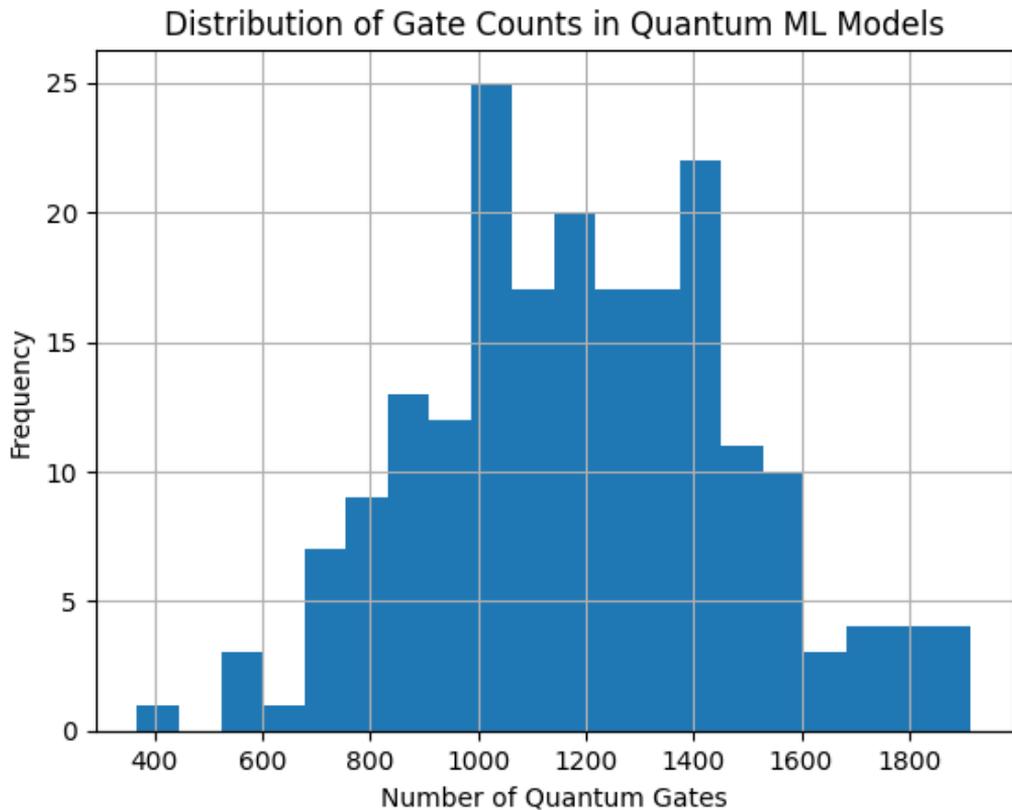


Fig 3: Distribution of Gate Counts Across QML Models

Classes of quantum circuit expressivity (Clifford circuits (efficiency of classical simulation) and IQP circuits up to universal quantum circuits) are used as a guideline to design architectures as they determine the minimal amount of quantum computational power necessary to solve specific learning problems. Recent algorithms Novel approaches to quantum machine learning models focus on the quantized cost of thing: measurement is stochastic and requires repeated measurements, gradients can have barren plateaus and quantum states can not be efficiently cloned to be evaluated in parallel. Extensional training variations of quantum algorithms Gradient-based optimization with the parameter-shift rule forms the base training approach in variational quantum: by using the forward pass of a quantum circuit at modified parameter values, one approximates the gradient of the circuit output. The parameter shift methods can give the precise gradient with no approximation error and can be done with only quantum circuit evaluation and classical post-processing but with multiple circuit evaluation per gradient evaluation-typically two circuit evaluations per parameter parameter shift rule. More advanced training Devices entail quantum natural gradient descent where gradients are preconditioned utilizing the quantum geometric tensor (Fubini-Study metric tensor) and thereby decrease on faster in specific optimization landscapes in trade off for gradient computation of the metric tensor. More efficient quantum natural gradient

methods that have been developed recently are the block-diagonal approximations to the metric-tensor, and random circuit-based stochastic approximation to the metric-tensor. Gradient-free optimization algorithms such as Nelder-Mead, genetic algorithms and Bayesian optimization have been used to solve quantum machine learning problems to eliminate barren plateau issues, but such algorithms generally need many more circuit executions to converge and may not scale to high dimensional parameter space. Methods of quantum-aided training employ quantum computers to implement itself, quantum gradient estimation algorithms and quantum linear solvers of second-order optimization, but again have challenges of complexity. The layer-wise training algorithms are a group of trainings of deep quantum circuits which are divided into sequentially trained layers, decreasing the barren plateau effect since they limit the circuit depth during training. Meta-learning methods learn through classical networks to optimize initializations of quantum circuit parameters, and thus can be used to learn with a few shots in which quantum circuits can be easily reconfigured to new tasks. Recent improved training algorithms are quantum imaginary time evolution of ground states of quantum Hamiltonians, counterdiabatic driving protocols to rapidly solve quantum optimization, and training with tensors training to optimize quantum circuits in the form of matrix product states. The training quantum machine learning models can be characterized by complexities in the number of circuit executions needed to achieve a given accuracy, as well as the classical cost of training parameter updates, gradient computation, and bookkeeping which is the source of complexity in near-term quantum circuits (limited number of qubits). Quantum kernel methods implement quantum feature maps but do not need to train quantum circuits, complicating counts of circuit executions needed to train quantum circuits. In kernel methods, inner products of quantum feature states in the quantum computers are computed using quantum kernels to implement quantum feature similarity between data points on quantum feature spaces. The advantage of quantum kernels methodology is based on its ability to compute quantum kernels more quickly than classical kernels, and the capability of quantum feature space to define relationships between data that are inefficiently defined by classical kernels. Quantum kernel methods are carried out in three steps, which are: preparing quantum-background state of the data points, estimating the entries of the quantum kernel matrix by quantum measurement, and applying classical kernel machine learning algorithms (support vector machines, kernel ridge regression) to the quantum kernel matrix. Quantum kernel methods Recount quantum kernels Tribal quantum kernels Reference frame quantum circuits involving parameters that are trained to maximize the quality of the kernel matrix, quantum Gaussian processes Risque quantum priors in a space of functions, quantum multiple kernel learning Trainable quantum circuits which fuse multiple quantum circuits to form a quantum kernel. Complexity analysis This can be shown to be a complex question: naive quantum kernel algorithms can scale the construction of the kernel matrix of n training examples with n^2 circuit evaluations, quadratic and prohibitive at large scale datasets. Sophisticated quantum kernel methods solve this by randomized quantum

circuit approximations which compute estimations of kernel functions with a fixed count of quantum circuit evaluations regardless of the amount of data, quantum cross-validation techniques that effectively estimate generalization efficiency without determining the entire kernel matrix in a single evaluation, and incremental learning techniques which, on evaluating each training instance in series, compute estimations of the entire kernel function without performing a complete quantum circuit evaluation. Complexity-theoretic tools have been used to define the expressivity of quantum kernels, with some quantum kernels (in particular those introduced by IQP circuits) also showing infeasibility to be had in estimating them by a classical computer under reasonable complexity assumptions, supporting a route to provable quantum advantage in quantum learning of kernels. Noise in practical quantum kernels would introduce error Alterations in the estimation of a kernel matrix that can negatively affect machine learning; these errors would be reduced by noise-resistant kernel matrix estimation methodologies such as kerber roducing quantum noise with classical post-processing techniques and quantum error mitigation methods. The main quantum mechanism models to generate effective models are quantum circuit Born machines, quantum generative adversarial networks (QGANs), and quantum Boltzmann machines that are characterized by different levels of complexity. Quantum circuit Born machines utilize parameterized quantum circuits to prepare quantum state where training is performed using maximum likelihood estimation by matching sample frequencies of quantum observation with target distribution frequencies. Born machine training is determined by circuit expressivity and the fidelity estimate-explicit fidelity computation some exponential number of measurements necessary to compute the fidelity of a circuit in a statistical sense, and approximate fidelity metering based on their classical shadows or as a compilation of classical strategies, grow the measurements used in a derivative manner. The service of quantum GANs has an implementation of a quantum generator circuit and a quantum or classical discriminator, trained in adversarial competition. Quantum discriminators QGANs are computationally unstable during training and mode collapse is also worse due to quantum measurement noise whereas quantum resources can be used to reduce the quantum cost of QGANs but quantum distributions are restricted by classical discriminators. More recent quantum generative modeling algorithms such as quantum Wasserstein GANs in which quantum circuits are trained to compute approximate Wasserstein distance and quantum diffusion models based on classical diffusion models which have set state-of-the-art generative modelling performance have been introduced. Quantum Boltzmann machines are generative models based on quantum thermal states, with quantum annealing or quantum approximate optimization used to benefit in training the Boltzmann machine. Measuring overhead attributes to the complexity bottleneck in quantum generative models Because many quantum measurements are required to stricthen samples drawn to the distribution which it is learning, to overcome the shot noise costs, and a gradients to supply training needs many further measurements, which makes the measurement overhead. Improved methods to

cut overheads on measurements also include importance sampling schemes which target measurements to high-probability configurations, adaptive measurement schemes which learn which measurements have most information, and the use of classical neural networks to learn classical approximations to quantum generative tasks which can do learning in data-limited problems with much less quantum resource. When quantum transfer learning is being implemented, it is typical to first train quantum circuits on tasks with plenty of data available on the source, freeze some circuit parameters and then fine-tune the remaining parameters with limited data on the target tasks. Generalisation The complexity benefit of transfer learning can arise when during the pre-training phase, beneficial features or structure is learned, which features in the target task and can be generalised. Quantum transfer learning has been applied in solving problems such as the prediction of molecular properties where circuits trained on small molecules can be used to predict properties of larger molecules, and image classification where quantum circuits trained on natural images can be transferred to medical imaging problems. Recent quantum transfer learning Responses Recent quantum transfer learning approaches involve meta-learning architectures, which train quantum circuits to learn how to run learning algorithms- quantum circuits which can rapidly adapt to new tasks with a small number of examples- as well as multitask quantum learning, which trains single quantum circuits to run multiple tasks related to a task jointly, sharing quantum circuit capacity among tasks. The analysis of complexity of quantum transfer learning looks at the trade-off between the pre-training cost (possibly high but amortized over a large number of target tasks) and fine-tuning cost (which must be low to actualise practical benefits). Negative transfer, in which pre-training causes regress, and not progress towards, target task performance is also to be defeated in quantum transfer learning as in classical transfer learning although quantum circuits might be especially vulnerable to barren plateaus in parameter-space. The main methods of strategies that reduce negative transfer are: proper choice of which circuit layers to freeze and which to fine-tune, regularization methods that spot large differences with the pre-trained parameters, and circuit architecture search to identify transferable quantum circuit blocks. Quantum ensemble approaches apply to the quantum case various methods of classical ensemble (bagging, boosting, stacking) approaches. Quantum ensembles can be complex due to the implementation of the ensemble members in parallel (this necessitates many quantum processors, or time-separated quantum computation), or implemented in sequence (this eliminates quantum parallelism). The ensemble members formed through quantum bagging are the based on the training of quantum circuits on bootstrap samples of training data in order to minimize variance but multiple quantum circuit training operations are necessary. Quantum boosting trains quantum circuits sequentially allowing the correction of errors of the previous ensemble members, which can probably increase the accuracy, but adaptive retraining can become costly on quantum computers. Quantum stacking applies quantum circuit in the base learners, and quantum or classical meta-learners in order to integrate predictions. Recent quantum

ensemble algorithms are quantum random forests which approximate quantum decision trees by combining many quantum decision trees, quantum dropout, which uses quantum gates in a randomized fashion to form implicit quantum assemblies, and quantum snapshot ensembles which process quantum circuit snapshots at random times during training. The complexity required to implement quantum ensembles should take into consideration additional measurement and classical aggregation computational complexity (quantum measurements are required on each ensemble member to produce prediction), and diminishing returns to increased ensemble size. In the advanced methods, quantum ensemble distillation can be trained to approximate an ensemble of quantum circuits, after which inference becomes cheaper than training an ensemble, and dynamic ensemble selection, which can employ classical learning to decide which members of an ensemble to stimulate on a particular test input to minimise the average quantum resource consumption.

3.4 Challenges in Quantum Machine Learning Complexity

The historical and contemporary problems of quantum machine learning have complex requirements in the context of basic physics, computational theory, and implementation and technological limitations [10-12]. Awareness of these issues will be the key to establishing a realistic timeline regarding quantum machine learning and also to focus the research work on the most effective problems. These obstacles include the basic complexity-theoretic constraints, to hardware flaws in the nearest future, to algorithmic complexity to systems integration. Barren plateaus are present when the graphs of quantum circuit cost functions are exponentially sinking and therefore the gradient-based optimization process fails. The underlying mechanism in the case of plateaus at beginning of quantum circuit is that, as the circuit depth and the number of qubits grows, quantum circuit gradients (quantum circuit gradients) are concentrated around the zero, which is due to the exponentially large quantum state space being explored by arbitrary parameterized quantum circuits randomly. Recent theoretical work has found that the phenomenon of barren plateaus occurs in a variety of situations global cost functions averaging the whole quantum system, quantum circuits initialized randomly, and hardware-efficient ansatzes with some symmetry. The implication of complexity is dire—not only can a quantum machine learning model have sufficient expressivity to model the functions that one would like to, but even when an appropriate set of parameter settings exists the training problem can be computationally infeasible because of exponentially decreasing gradients. Solutions to barren plateaus have been proposed in the form of local cost functions which limit optimisation to subsystems and avoid the global averaging that leads to vanishing exponential gradients; correlation-induced barren plateaus that also carefully design the circuit structure in order to avoid destructive interference of gradients; a proposal of initialising circuits with classical pre-

training or physics-based parameter values which initialises circuits close to good points before quantum optimisation; and layer-wise training, which breaks down deep circuits to interlocking superficial layers of stepwise trained circuits. But these mitigation techniques, too, have their concerns, such as local cost functions are not always sufficient to reflect targets of global optimization, careful circuit design can restrict expressivity, and layer-wise training makes more circuit calls. It is not yet clear that the efficient gradient training of deep quantum circuits can be done as this is a crucial problem of complexity-theoretic complexity in scale of quantum machine learning because quantum hardware noise and errors severely limit the scale and complexity of quantum circuits which can be reliably implemented on current and near-term quantum processors. The gate errors of single-qubit gates range between 0.1 to 1% and the errors of controlled 2 qubit gate are between 0.5 and 5%. Individual qubit gate errors accumulate quickly in quantum circuits. With a 1000 gate quantum circuit with a 1 per cent error per gate, the fidelity of such circuits is expected to be approximately equal to e^{-10} [?] 0.00005, so deep circuits are of no value to error mitigation or error correction. The first is noise in the gradient estimates used to train quantum circuits, which is usually 1 percent to 5 percent that corrupts quantum circuit outputs. The coherence time is limited by decoherence, which causes quantum properties to be lost due to the interaction with the environment; to a few milliseconds (for superconducting qubits) or milliseconds (for trapped ions), which restricts the port time of the circuit. The complexity issue is that novel quantum machine learning models, to achieve the same expressivity as classical neural networks, often need hundreds or thousands of gates of circuit depth, which is many orders of magnitude beyond anything achievable with state of the art hardware with useful fidelities. Quantum error mitigation has quantum error mitigation methods, such as zero-noise extrapolation, probabilistic error cancellation and symmetry verification, which can reduce effective error rates (below 2×10^{-2}) by classically post-processing quantum measurement outcome (with high measurement overhead), but at a high cost in classical computation. Error control in quantum computing Programs Making errors will no longer occur Quantum error correction uses logical qubits in the quantum algorithm to be encoded into physical qubits, and then error detection and correction achieved by means of fault tolerance is required. And making such a program fault-tolerant requires fault-tolerant quantum computers way beyond the current quantum computing capabilities. Most quantum error correction schemes such as surface codes need physical qubits on the order of 1000 per logical qubit to obtain more or less useful logical error rates at their current physical error rates, and introduce overhead requirements of 10-100x on the depth of circuits because of syndrome measurements and error correction. Complexity Implication This implies that the quantum machine learning that is fault-tolerant (capable of executing arbitrary depth quantum circuits reliably) will require quantum computers with millions of physical qubits and advanced error correction-resources which will not be attainable at least in 10 years. Water Tight quantum machine learning studies should thus be concerned with shallow circuits that

can run on noisy hardware and thus maybe a limiting factor to the quantum advantage that quantum machine learning can achieve on classical data. The data encoding bottleneck is a refined, but fundamental complexity problem of quantum machine learning on classical data. The majority of machine learning applications have as inputs classical data i.e. images, text, financial records, sensor measurements which need to be encoded into quantum states before quantum processing. Amplitude encoding (a non-linear encoding technique that provides exponential compression, i.e. N classical values are encoded into the amplitudes of $\log(N)$ qubits) is often cellularized with quantum circuits of depth $O(N)$, unless special structure arises. In the general machine learning pipeline, this complexity can cancel quantum speedups, i.e. in any quantum algorithm that has $O([\cdot]N)$ query complexity, anyway, any speedup of quantum to encode information must leave the overall complexity at $O(N)$ which is no better than classical methods. QRAM layouts have been suggested so that it can be loaded into depth $O(\log N)$ -data but the real-life QRAM implementation has been challenged immensely by the three challenges such as coherence time demands, hardware complexity scaling and the thermodynamic cost fundamentals. Other encoding methods do not encode the amplitudes, but angle encoding needs amplitude encoding, and basis encoding needs no amplitude encoding but only n feature representations with no compression, and has a limited expressiveness. Both spheres of quantum machine learning on classical data pose a complexity-theoretic problem: either highly-efficient data encoding algorithms need to be created, or quantum machine learning needs to be applied to problems with data that is inherently quantum (quantum sensor data, quantum simulation outputs), or quantum and classical hardware should be carefully synchronized to ensure that merely a minimal amount of data be transferred. Some more recent work has looked at strategies of lazy evaluation, in which quantum circuits are constructed in such a way as to load only that data on which a particular computation is to be performed but there are also uncertainties related to sample complexity in quantum machine learning on issues regarding whether quantum schemes are able to learn on fewer samples than classical schemes. Although quantum query complexity benefits are proven to be robust to particular problems, the sample complexity translation in learning systems is not clear. This difficulty is due to the fact that the complexities of samples are not only related to query complexity, but also on the loss function, sample noise, the complexity of hypothesis classes, and the output requirements. According to some theoretical findings, quantum learners need the same sample complexity as classical learners to PAC learn classical concept classes as the VC dimension (which is a measure of PAC sample complexity) is a property of hypothesis classes and not of computational model. Nonetheless, there are other outcomes demonstrating quantum sample complexity benefits to particular instances such as quantum PAC learning where both hypotheses and examples can be quantum, learning quantum states, and learning quantum statistical queries. The implication of practical complexity is unclear: although quantum algorithms can scale up their computational speed, it may not badly decrease the number of labeled examples needed

to support supervised learning, and so does not scale well to the applications where there is limited and expensive labelled data. Recent studies have examined quantum active and quantum semi-supervised learning to minimize the use of labeled examples, although these methods have their own complexity concerns such as the cost of labelling a choice of examples and the threat of selection bias. The most fundamental question of whether quantum computers can learn with exponentially fewer samples than classical computers with natural machine learning problems is still open and nature of the unresolved critical gap in quantum machine learning complexity theory. Such quantum algorithms as quantum principal component analysis or quantum recommendation system generate quantum states representing a solution, which should be measured and processed on a computer in the classical manner to obtain valuable information. Measurement processes should be able to measure the system being studied, and multiple measurements are necessary to approximate the expectation value with a given level of accuracy but this means that the process of measurement collapses quantum states and to approximate the expectation value to an accuracy of ϵ it requires $O(1/\epsilon^2)$ measurements. When a quantum algorithm has quantum state with exponential amount of information but the measurement can only extract the information that is of polynomial amount, the effective use might be restricted. The dequantization research program has given exposures that in certain issues in which quantum algorithms purported to be exponentially faster, on closer examination of the output needs it turns out that classical algorithms can be run to the same performance with almost comparable output, that is, by attaching similar quality to its classical output. To take the case of quantum algorithms that run recommendation systems and produce quantum states representing the recommendations, it may not be practically beneficial to use quantum instead of classical algorithms, considering the cost of extracting the recommendation list out of a quantum state through measurement. Its complexity challenge is to find quantum machine learning problems such that quantum outputs are either explicitly useful (in quantum applications such as quantum control or quantum simulation), can be measured efficiently to encode useful classical information, or that offers useful properties despite taking into consideration measurement overhead. Recent research on quantum machine learning with classical output has concentrated on either designing quantum algorithms that produce outputs which are themselves classical (by careful selection of observables measured) or problems whose quantum benefits remain despite measuring demands, but this remains a developing field of research with many outstanding questions. Scalability issues include both questions about the other way around, how complexity in quantum machine learning response to a given problem scales, and questions about how to scale implementing quantum machine learning algorithms to a larger quantum system. Theoretical scalability problems involve the critics of whether quantum machine learning algorithms can retain their benefits when dimensionality of data, the size of data, and the complexity of a model increase. There are quantum algorithms which have a provable execution speed only at a certain problem size, or in the asymptotic analysis

which may not be representative of real-life problem sizes. The existence of quadratic improvement of quantum speedups, including that of Grover-based amplitude amplification, might not be of practical value at all when constant factors and the presence of problem-specific structure are taken into account. Scalability issues that have been investigated in practice are limitations on the number of qubits (current quantum processors only support hundreds of qubits, which is a small fraction of the number needed to scale fault-tolerant machine learning), limitations in qubit connectivity (many quantum processors can only support nearest-neighbor qubit interactions, which is extremely inefficient in implementing fully-connected quantum neural networks), bottlenecks in classical controllability (generating control signals to scale to thousands of qubits at nanosecond timescales), and the need to support dilution refrigeration (millions of qubits are). The systems of quantum machine learning do not only increase with the number of qubits, but they also include the classical architecture to control, calibrate, and provide error-correction to quantum computers. More recent scalability studies have investigated modular quantum computing platforms in which smaller quantum processors can be linked together using quantum communication channels and used to perform distributed quantum machine learning, but in this case, this introduces the complexity of communication, and the types of entanglement accessible are limited. More scalable quantities of quantum computing Photonic quantum computing and topological quantum computing are also proposed to be a more scalable platform to qubit devices but these technologies are in turn themselves lacking in development, years and decades away. The absence of quantum machine learning theory is a fundamental challenge- there are many areas of quantum machine learning poorly represented fundamentally. Theoretical gaps A critical theory is required to address: Characterization: Which system of quantum machine learning models can be trained efficiently in the classical computational model and which cannot (whether quantum advantage can ever be realized); generalization bounds: To what extent can quantum learners be trained successfully on new data: when do quantum models trained on finite data model quantum novel data; expressivity Theory: What functions can quantum circuits possibly model, and how much efficiently, classical PAC, online, reinforcement learning models To what extent can quantum machine learning models be trained successfully on new data: when quantum models are trained on finite data will The lack of detail in theory renders it challenging to forecast which approaches in quantum machine learning will be hunched, to create quantum algorithms with assurances of performance, and discover lucrative areas of use. Recent theories have started filling these gaps: quantum learning theory of quantum data, generalization guarantees depending on quantum circuit complexity, expressivity analysis with network of tensors all that is known, still, is just the beginning. The complexity-theoretic difficulty is that quantum machine learning is in an exponentially large quantum state space, and exhaustive analysis of the quantum state is infeasible, and new mathematical tools and methods of proof are needed. The theory of quantum machine learning is an inherently

interdisciplinary field that only sustained research can develop; challenges to algorithm design in quantum machine learning include the need to figure out how to design quantum circuits that are both expressive (can model desired functions) and trainable (have no barren plateaus and pathological consistency problems) and implementable (meet hardware constraints). The three desiderata are often mutually exclusive more expressive circuits can be more susceptible to barren plateaus, hardware-efficient circuits can be less expressive and trainable circuits can need structures that are not compatible with hardware connectivity. The latest design work on algorithms has investigated structured ansatzes via problem symmetries, physics-inspired circuits via quantum simulation, and neural architecture search in order to find quantum circuits that optimize the expressiveness trainability implementability trade-off, although design principles are still lacking. The other challenge that is faced in the algorithms is the unavailability of quantum machine learning primitives- building blocks that can be re-used to make elaborate quantum learning systems. The advantages of classical machine learning are that it has rich primitives libraries such as activation functions, normalization layers, pooling operations, attention mechanisms, and loss functions, whose properties and complexity characteristics are well understood. Quantum machine learning does not yet have similar libraries of showed quantum primitives. The creation of quantum primitives involves getting to know what components of classical machine learning can be implemented efficiently using quantum algorithms, which can only be implemented using fundamentally new quantum algorithms, and which might be wasted in quantum learning systems. This is made worse by the fact that quantum circuits are required to be unitary (excepting measurement) and enables the implementation of classical non-linear operations and probabilistic components of classical deep learning.

3.5 Opportunities in Quantum Machine Learning Complexity

Even though the obstacles are very significant, quantum machine learning has an incredible potential of computational complexity breakthroughs, new applications, and scientific enhancement. Such opportunities are due to the special properties of quantum systems or the interaction of quantum computing with machine learning, as well as a blistering development of hardware and the speed of developing algorithms. It is important to balance realistic determination of existing limitations and visionary thinking of potential future opportunities in order to identify them and exploit them.

The nearest and possible quantum machine learning complexity opportunity is quantum advantage on particular problem structures. Instead of attempting to discover universal quantum speedups regardless of the problem faced by a machine learning application, studies are determining which problem properties and features can be better handled by quantum methods. The issues of quantum data, i.e. learning properties of quantum states,

predicting the dynamics of quantum systems, or optimising quantum device parameters, are more naturally addressed by quantum machine learning as the data is in a quantum form and one is no longer required to have an encoding bottleneck. An example of this opportunity is quantum chemistry machine learning to predict quantum simulation data, such as quantum circuits that learn the properties of materials by reading quantum simulation data, or pharmaceuticals by reading quantum simulation data, which further illustrates the rapid medical discovery of materials or drugs. High-dimensional optimization with targeted landscape properties is also a promising problem structure, where quantum tunneling and quantum annealing are potentially able to explore tricky optimization problem landscapes more effectively than quantum thermal annealing, especially in the case of tall narrow barriers between local minima. Recent work has found quantum advantage opportunities in quantum kernel methods to classification problems where data has a given geometric structure in quantum feature spaces, quantum principal component analysis of low-rank matrix problems, quantum Boltzmann sampling of particular probability distributions. Its complexity opportunity is to strictly characterize problem structures that are amenable to quantum advantage, such that leaders in the domain of algorithms can understand when quantum methods can work and all practitioners can prominently focus on when such methods will apply. To speed up quantum machine learning, a taxonomical approach to complexity of problems defining which structure of problems would be solved using quantum algorithms or classical algorithms would be created to prioritize resources on the best opportunities.

The potential offered by quantum computing and machine learning is synergistic, which opens more opportunities beyond the utilization of quantum computers to machine learning issues. Machine learning can be used to improve quantum computing by optimizing quantum circuits, training decoder quantum error correction, training quantum device calibration, and tuning quantum algorithm parameters. Classical machine learning has also already succeeded in optimization of quantum compilation (learning to compile efficient quantum circuits to implement target unitaries), in the design of quantum error correction decoders, and quantum devices calibration models, which have been shown to enhance gate fidelity transitions. The complexity opportunity is that this synergy can result in quantum computers being able to achieve useful performance at a rate lower than possible by improving hardware alone- machine learning optimization of quantum systems can successfully reduce error rates and increase circuit depths and improve quantum processor capabilities. On the other hand, quantum computing can help make entirely new machine learning accessible that can not be done classically. Quantum generative models which implicitly describe quantum correlations, quantum reinforcement learning agents which explore quantum state spaces, and quantum meta-learning systems which learn to learn quantum tasks are new models in machine learning made possible by quantum mechanics. The complexity-theoretic opportunity is that quantum enhanced machine learning can describe new

complexity classes and computational capacities not described by the classical learning theory and new vistas in computational learning may emerge.

NISQ-based quantum machine learning in the near-term offers potential prospects in improving the hardware development of quantum devices as well as the design of useful algorithms. With noise and qubit count-related restrictions, NISQ machines are considered good enough to investigate quantum machine learning in the hybrid between simulation at classical and fault-tolerant quantum computability. Variational quantum algorithms constitute a practical prospect of capitalizing on NISQ systems using the hybrid quantum-classical schemes of state preparation using quantum circuits and state training using classical circuits. Opportunity complexity The NISQ quantum machine learning complexities Opportunity NISQ quantum machine learning can identify specific tasks that require quantum-useful circuits on modest quantum devices, and are benefitably better than classical algorithms. In the recent past, industrial uses of NISQ quantum machine learning have been found in finance (optimizing a portfolio and analyzing risk), materials science (predicting material properties), and operations research (optimizing logistics), and results have shown that even quantum circuits with noise can help to improve the quality of a solution to some problem instances. The future of quantum computing that scale does not necessarily become universal quantum supremacy but tasks across problems generally with specific requirements where NISQ computers exhibit quantum (or classical) computational improvements. Moreover, NISQ machine learning motivates the hardware development through tangible benchmarking programs which indicate hardware weaknesses and inform the improvement criteria. ML with quantum machines is applied with loads inaccessible with any single simple quantum benchmark circuit, which uncarily reveal problems of performance in methods related to calibration, crosstalk, and coherence that can be solved by the architect of hardware. This mutual development of algorithms and hardware seizes a vicious nexus that hastens the process of quantum computing progress.

The quantum-inspired classical algorithm is a very intriguing potential that quantum machine learning studies open up. The algorithm design process of quantum machine learning has theolazed to inspiring new classical algorithms exploiting part of the quantum benefits even without quantum hardware. Initially used to simulate quantum many-body systems, tensor network methods have also been used to classify classical machine learning with impressive learning performance demonstrating, in some cases, state-of-the-art results, and theoretical guarantees of expressivity and generalization. The concept of dequantization research program is considered to be negative initially with regard to quantum machine learning, but has shown to have quicker classical algorithms with respect to problems like recommendation systems, low-rank matrix operations and even some supervised learning problems. The complexity opportunity is that quantum machine learning research expedites the efforts of trying classical algorithms by

compelling classical computational complexity to receive serious analysis and by offering novel algorithmic concepts based on quantum information theory. Quantum-inspired classical algorithms tend to have better scaling or other theoretical aspects, or other features never seen in classical algorithms, and in practice have advantages even without quantum hardware. The implications of this possibility are that quantum machine learning studies are providing value to quantum computers, and generally the field of algorithmic machine learning. The conceptual promise is to systematically explore quantum algorithms to gain classical understanding of such concepts as which quantum approaches have classical counterparts, and how we can know when quantum benefits are achievable and when they are not, and what known issues in quantum research have quantifiable traditionally resistant settled classical counterparts.

The potential applications of quantum machine learning to scientific discovery Scientific discovery entails opportunities to tackle some of the most difficult computational problems in science. Quantum simulation and quantum machine learning can be used to predict properties of quantum materials, design new catalysts of clean energy, and discoverment of high-temperature superconductors and modeling of intricate chemical reactions. The complexity advantage comes into place since the systems under consideration are quantum mechanical by their nature and classical simulation needs exponential resources whereas quantum one can be polynomial. Quantum machine learning builds on quantum simulation to learn useful representations of quantum states (quantum neural network state ansatz), forecast outcomes of simulation without quantum simulation (quantum transfer learning between smaller and larger system), and identifying physical information about quantum simulation data (quantum representation learning). Recent applications include quantum machine learning predicting significant quantities such as electronic structure of molecules with chemical precision accuracy, learning quantum phase transitions in condensed matter systems and finding novel materials to use in batteries and photovoltaism. It is possible to further extend to the fundamental physics where quantum machine learning can help in answering remaining questions in quantum many-body physics, quantum field theory and quantum gravity. The complexational-theoretical opportunity is that the discovery problems of scientific discovery tend to be natural quantum-structured problems, with embarrassment to compute resources undergoing such encoding paths, instead of directly using quantum factors of determination. The achievement of quantum machine learning in science would not only be of benefit in the advancement of particular areas of science but would prove quantum computing to be a revolution in science and would have the potential to accelerate investment and development.

Efficient quantum-classical optimized architectures of machine learning where the calculation is optimally divided between quantum and classical computers are a promising real-life opportunity in the near future. Instead of running complete machine

learning pipelines on quantum hardware or a combination of classical computing and quantum resources, hybrid models utilize quantum resources to address particular computational bottlenecks and use classical computing to address those areas where classical computing is better. Its complexity opportunity is to determine how best to divide labor between quantum and classical computing to apply to particular machine learning problems to use least computational resources and maximum accuracy. The recent hybrid designs are quantum feature extraction (with classical classification (using quantum circuits to compute features predicted by classical models) and classical preprocessing (with quantum optimization) and quantum sub routines (in classical algorithms) (using quantum circuits to perform certain operations in a classical learning algorithm). Hybrid architecture design needs to take into consideration classical-quantum communication overhead, quantum measurement costs and coordinating distributed quantum-classical computation complexity. Further hybrid approaches have been advanced such as variational quantum-classical networks/layers of neural networks that alternate quantum and classical layers, co-optimization of quantum-classical networks/layers creating a single quantum-classical model and quantum-classical ensemble approaches that combine quantum and classical predictors. The opportunity of theory is to create complexity theory of hybrid quantum-classical computation defining when hybrid strategies are significantly beyond either a pure quantum or pure classical strategy, with rigorous foundations on hybrid system design.

Quantum machine learning services and workstations offer potentials to make quantum computing more democratic, and also to speed up the time required to develop algorithms. Full-fleet quantum machine learning models which hide the hardware implementation allowing easy to use applications allow domain scientists and machine learning users to explore quantum techniques without background in quantum physics. The tools such as PennyLane, TensorFlow Quantum and machine learning Qiskit include the ability to build quantum circuits, automatic differentiation and interoperability with classical machine learning libraries. The complexity opportunity is that improved tools will take shorter time to research because they will help cut down on the overhead of implementing quantum machine learning concepts. The use of circuit compilation tools that automatically compile quantum circuits to particular hardware, quantum-classical workflow systems offering management of hybrid computation, and quantum cloud systems offering on-demand access to quantum processors make quantum machine learning research less accessible to new entrants. Ecosystems develop based on benchmarking platforms with strict comparisons between quantum and classical methods on standardized problems, quantum algorithm catalogs databases with known quantum strengths and weaknesses, and quantum machine learning learning platforms with technologies that educate the future quantum scientists. This can be extended to specialized quantum machine learning accelerators, i.e. custom quantum hardware designed to execute machine learning programs and not general quantum

computation, which can potentially have a better performance-cost trade-off than universal quantum computers on particular machine learning programs. Quantum machine learning platforms construction is a software innovation opportunity that can provide value regardless of hardware advancements of quantum hardware, allowing scientists to be ready to face the quantum machine learning future even as quantum hardware continues to improve.

The opportunities of quantum-enhanced data privacy and security in machine learning are unique capabilities that come up due to quantum information-theoretic security. The quantum machine learning protocols are capable of making privacy guarantees that are classically impossible by quantum cryptography, quantum homomorphic encryption, and quantum secure multi-party computation.

3.6 Impact of Quantum Machine Learning Complexity Research

Computational complexity of quantum machine learning model research has an influence in many dimensions, including theoretical computer science, developing practical algorithms, hardware design, industry, and socially significant aspects of artificial intelligence. These impacts offer insight into why quantum machine learning complexity studies have relevance beyond the academic level of interest and influence on the future of both directions of quantum computing and machine learning.

Research in quantum machine learning complexity has significantly influenced computational theory of the quantum perspective in other fields by providing new understanding of the complexity classes in quantum computation, learning theory, and the inherent limits of quantum computation. Studies of the complexity of quantum learning have resulted in improved definitions of complexity classes such as BQP, identification of oracle separations between quantum and classical learning and quantum analogs to classical learning models such as PAC learning, online learning and statistical query learning. The study of quantum query complexity in learning problems has found some subtle connections between various measures of complexity query complexity, sample complexity, and communication complexity, which have developed the general complexity theory. Though originally inspired by quantum machine learning, the dequantization program has led to new classical algorithmic methods and tools of complexity theory such as improved sampling methods and matrix concentration bounds. The classical complexity theory has also been pointed out as an area of gap in quantum machine learning complexity research, where classical complexity is poorly understood and classical computational limits find reason to be studied. This has been applied to quantum information theory by characterizing quantum entanglement as computational resource in learning, studying quantum channel capacities to learn problems and exploring quantum communication complexity in distributed learning.

Such theoretical developments offer the rigorous grounds which are required to make plausible arguments on the benefits and limitations on quantum machine learning and avoid either over-confidence or unreliable cynicism.

It has had a significant practical effect on the development of quantum algorithms, such as complexity analysis to help guide the design of more practical quantum machine learning algorithms and recap computation costs in proposed algorithms. Strict complexity analysis has resulted in better quantum algorithms finding bottlenecks to optimise, finding better quantum circuit compilation methods, and hybrid quantum-classical methods that reduce the number of computeable resources overall, and not just optimise quantum or classical elements. Complexity research has found that a great deal of early quantum machine learning work had made an unjustified high claim of benefits by ignoring the cost of encoding data, the overhead of measurement, or the need to produce output, and has made more earnest and careful claims about algorithms in more recent work. The complexity analysis of the barren plateau phenomenon has radically changed quantum machine learning algorithm design that stimulated local cost functions, problem-specific ansatzes, and alternative training schemes. Complexity-inspired algorithm design has led to useful contributions such as quantum kernel algorithms, vQAs with convergence guarantees and quantum-classical hybrid algorithms with strength-of-numericity guarantees. The effect is also felt in terms of measuring the complexity of the algorithm benchmarking, where the complexity analysis may make theoretical projections that can be used to guide experimental validation work and in interpreting quantum hardware results. The idea of complexity research itself has been the catalyst behind the design principles of quantum algorithms, or more generally, said principles of constructing quantum circuits with certain desired complexity properties, which have spurred the design of new quantum learning algorithms much faster.

The research on quantum machine learning complexity has strongly impacted the development of quantum hardware, by defining hardware needs of practical quantum learning and driving particular hardware development. Complexity analysis shows that to realize the benefits of quantum machine learning, quantum processors with capability areas such as: adequate number of qubits to scale with problem size, low error rates to run problems with lots of circuit depth, fast gate execution to run long enough, and connectivity to support ability to run its problems on an efficient circuit are needed. They are based on complexity analysis to work on these requirements and they offer specific targets to hardware engineers. The observation that closer-term applications of quantum machines need just scaleable numbers of qubits yet the higher-level gates must have better fidelity has motivated the creation of hardware to focus on error reduction, in addition to qubit scale. Specialized hardware features such as mid-circuit measurement in quantum feedback, qubit reset in circuit reuse and fast classical-quantum communication in hybrid algorithms have been prompted by complexity analysis on

particular quantum machine learning workloads. It has enabled the hard-to-implement quantum machine learning algorithms to be identified to scale to the pattern of hardware connectivity, which has resulted in a virtuous cycle of algorithms being reformulated to limits hardware can implement, and hardware being developed to implement significant algorithms. Based on the complexities analysis, quantum machine learning benchmarking yields standardized tests of quantum processor performance beyond the microbenchmark, which can be used by hardware developers to evaluate the capability of real-world quantum systems. The reduction is the ability to inform qubit technology choice, to different quantum computing platforms (superconducting qubits, trapped ions, neutral atoms, photonic qubits, topological qubits) having different complexity trade-offs (gate speed, error rates, connectivity, etc.) as well as quantum machine learning complexity needs aiding in the assessment of which technologies are most appropriate to which tasks.

Complexity research has influenced the way quantum machine learning has been adopted in industry by identifying realistic near-term uses, and controlling the expectations of quantum benefits. Several companies, such as IBM, Google, Amazon, Microsoft and many startups, have invested in quantum machine learning heavily on the basis of potential value based on complexity-theoretic analysis being conducted on the value. Complexity research gives industrial investing a sense of what quantum machine learning applications will get to near-practical benefit soon (quantum chemistry, materials science, optimization of a particular problem structure) and what are years off (general quantum neural networks, quantum deep learning) based on fault-tolerant quantum computers. The sincere review of complexity constraints has somewhat worked against the hype by making more resources to concentrate on the realizable objectives. There has been an industry-academia collaboration on the complexity of quantum machine learning research with companies making quantum hardware accessible and supplying real-world problems and academics contributing complexity analysis and algorithm development to these problems. Other targeted industrial applications are the use of variational quantum algorithms to simulate chemistry by pharmaceutical and materials firms, quantum optimization to solve logistics and scheduling problems in operations-oriented firms, and quantum machine learning to model financial problems in banks and investment companies. Complexity based design of hybrid quantum-classical machines has especially influenced industry, furnishing viable avenues to derive utility out of existing quantum hardware without having to await fault tolerant quantum computers. Even applications of industrial quantum machine learning, where quantum advantage is not yet realized, would facilitate algorithm refinement and hardware improvement by enabling real-world testing, establishing feedback loops and accelerating the technology development.

Education also has educational influence on quantum computing education, machine learning education and interdisciplinary education programs to prepare quantum machine learning researchers. The quantum machine learning complexity-theoretic roots have also been used as pedagogically useful examples to provide conceptual understanding regarding important concepts in quantum computing by making abstract quantum concepts concrete through machine learning applications which are familiar to learners. The studies of quantum machine learning are currently found in computer science, physics, and engineering programs at universities around the globe, preparing quantum computing industry employees. The complexity of quantum machine learning that cuts across the disciplines of quantum physics, computer science, mathematics, and statistics leads to educational initiatives that does not reinforce disciplinary barriers, but trains students as interdisciplinary to work across a set of fields concomitantly. Education Access outsourcing Clicking courses, tutorials and open source software to quantum machine learning have made quantum computing education more accessible to students and professionals who do not have access to specialized quantum computing education programs. The effect has contributed to the creation of quantum machine learning learning applications, such as quantum circuit simulators to visualize quantum states and gates, interactive guides to teaching quantum algorithms by experimenting with a particular circuit, and specialized quantum machine learning learning benchmark datasets. Research in education about the ways students learn quantum machine learning ideas has found out the key misunderstandings and effective methods of teaching, which increases quantum computing education in general. Quebecbecpendance of quantum machine learning researchers can be added to the wider quantum computing industry workforce development project since, regardless of the application of quantum computing, skills in quantum algorithm design, quantum error analysis, and hybrid quantum-classical programming skills are transferable.

Quantum machine learning complexity research has fuelled scientific methodology with novel methods of solving scientific problems of possible computation, and strict algorithms to measure the amount of computation done. The rigor standards of systematic complexity analysis used in quantum machine learning have formed the standards of rigor in quantum computing research more broadly, by highlighting the need to specify the input models, output constraints and realistic resource accounting. Quantum computing scientific practice in developing benchmarking methodologies that sufficiently compare quantum and classical methods has progressed to overcome publication bias on positive outcome studies and promote truthful reporting of limitations. Research on quantum machine learning complexity has also contributed to scientific computing methodology thereby showing the usefulness of special purpose quantum-classical computing in solving computational problems specially with specialized hardware implemented-an approach that is also relevant to non-machine learning scientific simulation and data analysis. Combining machine learning with

quantum simulation in materials science, chemistry, and physics has led to emergent scientific paradigms, in which machine learning and computations are, respectively, intimately connected and coupled to efficient searching of scientific search spaces and to the objective and accurate assessment of candidate solutions. Such effects on methodology also apply to experimental science, where quantum machine learning can analyze quantum experimental data in real-time, design experiment adaptively to optimize experimental parameters and extract patterns in high-dimensional data of science.

The social dimension of the complexity research on quantum machine learning is only beginning to accumulate; therefore, the societal and economic dimension of this field is high in the form of the effects on economic competitiveness, technological sovereignty, workforce development, and responsible development of the powerful AI systems. Countries around the globe have launched quantum computing research efforts with the quantum machine learning promise being one of the motivations, in addition to quantum computing leadership being seen as strategic necessity in future technology competitiveness. The complexity minded realization that quantum machine learning is long term financially demanding, not a technology to offer immediate revolutionary demands, has aided in the realization of realistic approaches of funding research in the usual realms of cloudy fundamentals alongside applied development. Quantum machine learning casts its challenge on AI safety and AI health, quantum machine learning systems can have new types of failure, adversarial sensitivity, or unexpected behavior that will not manifest in classical models and AI safety studies will need to take novel approaches. Possibility of quantum computers increasing the capacity of machine learning also attracts concerns on technological inequality, given that access to quantum computing resources can end up being controlled by the rich countries and huge corporations, thus increasing the existing technological disparity. On the other hand the, access to quantum computing resources and open-source quantum machine learning software are more or less democratized through clouds with the result that more people can participate in research and application of quantum machine learning. Such complexity research showing that the benefits of quantum machine learning are problem-specific and no longer general gives some hope that quantum computing will complement and, not displace, classical machine learning, possibly, facilitating transitioning to better economic states as quantum technology advances. As ethical aspects, it should be ensured that quantum machine learning privacy guarantees are made available to vulnerable groups, that quantum-enhanced machine learning is not used maliciously by disinformation or surveillance, and that there is an open dialogue about quantum machine learning capabilities and limitations to reduce hype cycles which often result in loss of trust in science.

3.7 Future Directions in Quantum Machine Learning Complexity

The future perspective of quantum machine learning complexity studies will be comprised of theoretical issues that will shape out the basis of the field, algorithm issues that will allow practice, hardware issues that will extend the reach of quantum computing, and interdisciplinary issues that will work towards an expansion of the influences of quantum machine learning [7,13-16]. It is necessary to balance between short-term feasibility with important by analyzing the opportunities of long-term transformations, in understanding that research seeds of the future are in the current core research.

Strict definition of quantum-classical learning separation embodies a fundamental theoretical direction of the future which will dictate whether quantum machine learning can be brought to its full potential. The existing knowledge of the time when quantum learners are performing better than classical learners is still inconclusive where the empirical evidence demonstrating quantum advantage on certain problems exists together with the negative evidence that quantum advantage does not exist on certain problems. Future studies should mathematically define the threshold between quantum-advanced learning and classically-equivalent learning, finding the properties of the problem, structural constraints, and measures of complexity defining in which direction of the boundary particular learning tasks should be. In this direction, this research needs to develop quantum learning theory by developing quantum VC dimension and quantum Rademacher complexity among other quantum analogs to classical learning complexity measures and other new quantum-specific learning measures that can be more indicative of quantum learning power. The description must go further to the average-case complexity and instance-specific complexity with the realization that quantum advantages can be seen to instances of typical problems with worst-case complexity being comparable to classical complexity. The role of the quantum entanglement, quantum coherence and other quantum resources in the learning benefits should also be addressed in the future to find out whether and how the quantum properties of quantum space result in computational power. It will require the tools of quantum information theory, representation theory, algebraic complexity theory, and mathematical physics of advanced level, indicating that interdisciplinary effort will be necessary in the long term. Quantum-classical separation in learning solves would offer the theoretical framework that is required to forecast the quantum-accommodating learning issues to real-life learning challenges that may be served efficiently, as a resource well-being to study and cultivate quantum machine learning.

Training on post-barren plateau methods are a future essential algorithmic direction that will either enable interactions to be trained or not in deep expressive quantum circuits. Although the existing knowledge demonstrates the existence of the generic variational quantum circuits with known plateaus, the future studies need to create the systems of

training that are capable of avoiding or overcoming the plateaus without compromising on the quantum circuit expressivity. The promising directions are to create problem-specific problem-based initialisation protocols, based on classical machine learning, physics-based ground states, or optimisation using a tensor network that put parameters into favourable positions before adopting a quantum optimisation process; designing quantum circuit architectures with inductive biases informed by structure of the problem that avoid barren plateaus without inducing barren plateaus, but retain expressivity; exploring alternative optimisation processes based on alternative formulations of the cost-function, measurement, and definition of objectives that do not obey barren plateau behaviour; and higher-order optimisation algorithms. Other entirely new training paradigms not based on gradient-based optimization that may be researched in the future are evolutionary training of quantum circuits, quantum-based training of quantum circuits where training entails the execution of a task, and quantum-inspired classical training algorithms to explore efficient training strategies, trained on quantum circuits classically. No big algorithmic break even quantum machine learning applications will be possible without provably trainable quantum circuits with a sufficient expressiveness, which could potentially unlock the full capabilities of quantum neural networks and break out of barren plateaus. The proposed direction research would entail integrating knowledge in classical optimization theory with quantum control theory, quantum adiabatic computation theory, and numerical analysis, to come up with in depth knowledge of quantum circuit optimization landscapes and effective training algorithms which navigate these landscapes efficiently.

The ability to implement quantum machine learning with fault-tolerant quantum computers will be a direction in the long term that will determine the ultimate capabilities of quantum learning with hardware limitations eliminated. Although near-term research is both on NISQ devices with intolerable noise constraints, fault-tolerant quantum computers with error-corrected logical qubits will be developed in the future to support deep quantum circuits that are impossible at present. Future work should get ready with fault-tolerant quantum machine learning it should develop algorithms with fault-tolerant execution and study the resources requirement in logical qubits and number of T-gate and distinguish applications where the fault-tolerant quantum machine learning constitutes a transformative advantage. Fault-tolerant quantum machine learning represents some of the possibilities such as quantum deep learning with thousands of quantum layers of hierarchical quantum feature learning, quantum recurrent neural networks with long memory lengths to model complex temporal behavior and quantum generative models, which prepare high-fidelity quantum states to be used in quantum simulation and quantum sensing applications. Analysis of fault-tolerant quantum machine learning needs to take into consideration energy consumption of errors correction overhead-quantum error correction uses a lot of resources even when errors are being successfully fixed- so that algorithms with quantum benefits large enough to

reimburse error fixing expenses are required. The research conducted in the future ought to consider the best algorithms of quantum machine learning given fault-tolerant quantum computers, and define which machine learning architectures, training and applications are the most promising in the fault-tolerant regime. The co-design of quantum error correction schemes and quantum machine learning algorithms are also considered as this direction of research as well as implementing specialized error correction schemes in response to machine learning workloads, and quantum circuits implemented to minimize error correction overhead. The roadmap to fault-tolerant quantum machine learning has intermediate targets such as early small-scale demonstrations of error correction, logical qubit systems with 10-100 logical qubits suitable to implement proof-of-concept quantum learning applications, and medium-scale fault-tolerant processors with 1000 or more logical qubits and which realize realistic quantum learning applications.

Table 1: Quantum Machine Learning Applications and Complexity Characteristics

Sr. No.	Application Domain	Primary Techniques	Complexity Advantage	Key Challenges	Implementation Methods	Future Directions
1	Quantum Chemistry	VQE, QAOA, Quantum simulation	Exponential for quantum systems	Barren plateaus, circuit depth	Variational algorithms, hybrid quantum-classical	Fault-tolerant quantum chemistry, multi-scale modeling
2	Drug Discovery	Quantum molecular docking, QNN property prediction	Polynomial to exponential depending on task	Sample complexity, measurement overhead	Quantum kernels, transfer learning	AI-guided quantum drug design, personalized medicine
3	Materials Science	Quantum generative models, inverse design	Exponential for property simulation	NISQ limitations, validation costs	Quantum Boltzmann machines, quantum optimization	High-throughput quantum materials screening
4	Financial Optimization	QAOA, quantum annealing	Polynomial speedup for certain QUBO	Data encoding bottleneck	Hybrid optimization, quantum sampling	Quantum risk analysis, real-time portfolio management
5	Natural Language Processing	QNNP, quantum transformers	Representational capacity advantage	Barren plateaus, deep circuits	Compositional quantum circuits	Large quantum language models,

6	Computer Vision	QCNN, quantum feature extraction	Theoretical exponential for certain transforms	Classical data encoding	Quantum convolution, hybrid architectures	quantum semantic analysis Quantum video processing, quantum 3D vision
7	Anomaly Detection	Quantum clustering, quantum autoencoders	High-dimensional search advantages	Real-time latency requirements	Quantum one-class SVM, quantum statistical tests	Quantum streaming anomaly detection
8	Time Series Forecasting	Quantum RNN, quantum reservoir computing	Temporal correlation capture	Short coherence times	Quantum-classical hybrid recurrence	Long-term quantum temporal modeling
9	Optimization Problems	Quantum optimization algorithms	Problem-specific polynomial to exponential	Approximation ratio limitations	QAOA, quantum gradient descent	Quantum metaheuristics, quantum constraint solving
10	Quantum State Learning	Quantum state tomography, shadow tomography	Exponential proven advantage	Measurement complexity	Adaptive measurements, compressed sensing	Scalable quantum state reconstruction
11	Reinforcement Learning	Quantum policy gradient, quantum Q-learning	Sample efficiency potential	Exploration-exploitation balance	Quantum advantage actor-critic	Quantum model-based RL, quantum multi-agent systems
12	Graph Machine Learning	Quantum walk algorithms, quantum GNN	Quantum walk mixing advantage	Graph encoding complexity	Continuous-time quantum walks	Large-scale quantum graph networks
13	Generative Modeling	Quantum GANs, Quantum Born machines	Distribution expressivity	Training instability, mode collapse	Quantum circuits as generators	High-fidelity quantum generative models
14	Federated Learning	Quantum secure	Information-theoretic privacy	Communication overhead	Quantum secret sharing	Private quantum

		aggregation				distributed learning
15	Meta-Learning	Quantum meta-learners	Few-shot learning efficiency	Transfer learning validation	Quantum MAML adaptations	Universal quantum learners
16	Active Learning	Quantum query selection	Sample complexity reduction	Query strategy complexity	Quantum uncertainty sampling	Quantum Bayesian optimization
17	Causal Inference	Quantum structural equation models	Confounder representation	Causal structure discovery	Quantum do-calculus	Quantum counterfactual reasoning
18	Topological Data Analysis	Quantum persistent homology	Exponential feature space	Quantum topology algorithms	Quantum Betti number computation	Quantum shape analysis
19	Probabilistic Programming	Quantum probabilistic circuits	Sampling efficiency	Probabilistic inference complexity	Quantum approximate inference	Quantum Bayesian networks
20	Continual Learning	Quantum memory consolidation	Catastrophic forgetting mitigation	Quantum memory requirements	Quantum experience replay	Lifelong quantum learning systems

Quantum-enhanced interpretability and explainability is an average future trend that satisfies the very urgent situation of how quantum machine learning models learn and why they predict in a particular way. Systems like quantum machine learning transforming the field of theoretical study into actual application need interpretability tools to enable debugging models, promote fairness and safety, and establish trust in systems. Future studies need to come up with quantum versions of classical interpretability tools such as feature importance analysis, saliency maps, attention visualization, and concept activation vectors, as well as investigate fundamentally quantum interpretability tools, such as quantum state tomography, quantum entanglement structure, and quantum circuit analysis. The difficulty is that quantum models act in exponentially large Hilbert spaces which quantum correlations do not have classical analogs, and may need entirely new interpretability frameworks. Future prospects of research include theoretical imaging of visualizing quantum circuit decision boundaries based on quantum state space analysis, understanding quantum features learnt by quantum circuits using quantum circuit surgery to isolate circuit components, and theoretical imaging quantum counterfactual explanation, describing how quantum inputs would have to be different to change the prediction. The implications of complexity are large-interpretability techniques that need exponential classical resources to interpret quantum models can be infeasible, and interpretability techniques that use quantum computation to explain quantum models can be devised efficiently. Also in

future research, the intrinsic interpretability of quantum machine learning models over the classical models should be researched in order to establish whether quantum models are inherently more or less interpretable than classical neural networks. Such a direction of research is related to quantum information theory by the question of what properties of quantum models can be effectively observed and to human computer interaction by the question of how to describe the explanations of quantum models to human consumers who do not have a background in quantum physics.

The domain-specific quantum machine learning architectures are a direction of the future, specialized quantum circuits based on a particular area of application, and may perform more directly than general quantum neural networks. Instead of building more universal quantum machine learning models, future work should build domain-specific quantum architectures that exploit properties such as molecular symmetries, electronic structure properties, risk correlations, market dynamics, grammatical structure, semantic compositionality, and visual hierarchy quantum convolution and pooling operations, respectively. The domain-specific architecture complexity advantage comes with the fact that domain knowledge is directly encoded in quantum circuit structure, search space per training is reduced and barren plateaus that afflict generic circuits are possibly avoided. The way in which classical machine learning domain-specific architectures (convolutional networks on images, transformers on sequences, graph neural networks on relational data and so on) can be represented in quantum contexts need to be examined (systematically) in future research to infer which classical architectural principles can also be efficiently implemented in quantum environments and which ones necessitate fundamentally new quantum design methodologies. Others found in this line of research are the creation of automated domain-specific search strategies of quantum architecture implementation, which identify optimal quantum circuit designs to particular domains, and the generation of libraries of reconfigurable quantum circuit units optimized to particular domain-specific tasks. The effect would be a practical quantum machine learning systems, both in respect to exploiting quantum computational power and domain-specific structure in realization of quantum benefits to the practical applications.

The security and robustness of quantum machine learning is a direction critical to the future because quantum learning system shifts toward usage in security sensitive systems. Future studies need to deal with various security aspects such as adversarial robustness of quantum models to adversarial examples that are constructed to cause one to misclassify, backdoor attacks where adversaries introduce malicious code in the training process of quantum circuits, model extraction attacks where adversaries seek to steal quantum models by use of queries and privacy attacks where adversaries can extract sensitive training information based on output of quantum models. Quantum facility comes with new adversarial issues and opportunities: quantum adversarial instances can use quantum features, such as entanglement, to be more effective than classical

adversarial instances, whereas quantum systems can also be constructed with a collection of fundamentally secure protocols that are impossible in a classical context. Future studies need to come up with certified robustness mechanisms of quantum learners that give verifiable assurances on the robustness of models, quantum-secure training schemes that deterring backdoor and poisoning threats on quantum cryptography and quantum privacy-preserving learning schemes on the basis of quantum secure multiparty computation. The analysis of the complexity of quantum machine learning security should take into consideration both classical and quantum adversaries, that is adversaries with classical computational power, and adversaries with quantum computers, to know which security properties are definable with what capabilities of adversaries. This direction of research can be related to post-quantum cryptography and quantum cryptography because both achieving security against a quantum adversary of classical machine learning and against a classical adversary of quantum machine learning necessitate crypto-foundations. The security implications of quantum machine learning studies on the society are huge as more long-distance applications of machine learning such as healthcare, finance, and autonomous systems where the security risk may be severely harmful are being deployed.

Cross-fertilization with classical machine learning is an endeavour with significant future prospects with the recognition that classical machine learning and quantum machine learning can benefit one another. Substantial scale works in the areas of quantum extensions of the advancements in classical deep learning need to be translated to quantum implementations, and quantum analogues of future directions in classical deep learning ought to be developed (quantum versions of classical and reinforcement-based neural networks, quantum training and testing of quantum computers), including quantum-based systems generating language and applying it to computational tasks. This translation should be two ways, as quantum machine learning models will catalyze new classical methods such as quantum-inspired attention models, quantum circuit-based architecture of a classical circuit, and quantum optimization-inspired optimization methods. Next generation quantum-classical architectures based on architectures that firmly couple quantum and classical learning should also be studied where quantum features extraction or quantum optimization subroutines are substituted with much more deeply hybrid quantum and classical architectures where quantum and classical components are co-evolved with each others training. The quantum-classical integration analysis should be complex enough to maximize the end-to-end performance based on quantum-classical communication cost, quantum-measurement overhead, and classical preprocess requirements. Such a future direction of the field involves tight integration between the areas of quantum computing and classical machine learning, disaggregating the divide which has occasionally characterized each of the two areas and making it possible to transfer the knowledge either way. It would have a speed up effect on both

quantum and classical machine learning, where quantum concepts would improve classical methods, and classical methods would improve quantum implementations.

The standardization and quantum machine learning benchmarking is one of the most significant future steps in the evaluation and equal conditions of quantum and classical methodologies. Future directions should include the establishment of complete quantum machine learning benchmark suites spanning a variety of problems, data types and regimes, and have standard evaluation procedures that provide reproducibility and a fair comparison. Both quantum advantaged problems where the quantum methodology is supposed to shine and general machine learning problems where the classical methodology reign supreme should be used as benchmarks to provide a balanced evaluation of what the quantum can do. The future work could standardize the complexity of quantum machine learning, which include end-to-end time complexity (including classical preprocessing, quantum execution, and classical post-processing), resource complexity (quantum qubit constraints, number of gates, gate depths), and energy complexity (total energy expenditure both classical and quantum). The performance prediction of quantum machine learning models that would predicts performance of quantum algorithms on a particular hardware, but are not fully implemented could improve the speed of developing algorithms by providing quick feedback on approaches under consideration. Related standardization efforts ought to be made on quantum machine learning software interfaces, quantum circuit intermediate representations, and quantum-classical communication protocols, to allow the compatibility of quantum machine learning models across quantum hardware platforms and the scalability of quantum learning algorithms compared to quantum hardware platforms. Proper attention to complex-theoretical basics of fair benchmarking includes proper attention to input models (how are benchmark datasets given), output demands (what accuracy or quality criteria are admissible), and comparison on the baseline (which classical methods are the state-of-the-art). Standardization of benchmarking could be done with the help of international cooperation via such organizations as IEEE, ACM, and quantum computing industry consortia, which guarantees the large-scale adoption in the community.

Table 2: Quantum Machine Learning Complexity Analysis Framework

S r. N o.	Complexity Measure	Classical Bound	Quantum Bound	Separation Result	Analysis Technique	Open Problems	Research Priority
1	Query Complexity	$O(N)$ for search	$O(\sqrt{N})$ Grover- optimal	Provable quadratic separation	Polynomial method, adversary bounds	Quantum learning query gaps	High - fundamental limits

2	Circuit Depth	N/A (classical circuits)	$O(\text{poly}(n))$ for BQP	Circuit-dependent	Circuit synthesis, gate decomposition	Optimal circuit depth bounds	High - NISQ constraints
3	Sample Complexity	$O(d/\epsilon^2)$ PAC bound	Quantum PAC unclear	Problem-specific	VC dimension, Rademacher complexity	Quantum-classical sample gaps	Critical - practical impact
4	Gate Count	N/A	Hardware-dependent	Architecture-specific	Gate compilation, circuit optimization	Minimal gate implementations	Medium - resource estimation
5	T-Depth	N/A	Fault-tolerance metric	T-gate-dependent	Clifford+T decomposition	Efficient T-depth reduction	Medium - future hardware
6	Measurement Complexity	N/A	$O(1/\epsilon^2)$ shot noise	Fundamental quantum limit	Statistical estimation theory	Measurement-efficient learning	High - current bottleneck
7	Entanglement Requirements	N/A	Task-dependent	Resource theory bounds	Quantum resource theories	Minimal entanglement for advantage	High - fundamental resources
8	Communication Complexity	Classical CC bounds	Quantum CC potentially lower	Exponential for some problems	Quantum communication protocols	Distributed quantum learning CC	Medium - distributed systems
9	Memory Complexity	$O(\text{model size})$	$O(\log N)$ for some tasks	Amplitude encoding advantage	Quantum RAM analysis	Practical QRAM implementation	Low - distant future
10	Time Complexity	Classical algorithm-specific	BQP containment	Conditional on $BQP \neq BPP$	Complexity class analysis	BQP vs BPP for learning	Critical - foundational
11	Approximation Error	ϵ -accuracy bounds	ϵ -accuracy quantum	Algorithm-dependent	Error analysis, concentration	Quantum approximation guarantees	High - quality assurance
12	Gradient Variance	Bounded by loss	Exponentially	Architecture-	Variance calculations	Barren plateau	Critical -

		landscape	vanishing (barren)	dependent		characterization	trainability
13	Convergence Rate	$O(1/T)$ typical	Quantum natural gradient faster	Problem-dependent	Optimization theory	Quantum optimization landscapes	High-training efficiency
14	Expressivity	Universal approximation	Universal quantum	Exponential Hilbert space	Quantum approximation theory	Efficient quantum expressivity	High-model capacity
15	Generalization Bound	PAC-Bayes, uniform convergence	Quantum generalization unclear	No proven separation	Statistical learning theory	Quantum capacity measures	Critical-deployment reliability
16	Energy Complexity	Joules per operation	Thermodynamic quantum limits	Landauer limit considerations	Physical energy analysis	Energy-efficient quantum learning	Medium-sustainability
17	Latency	Classical inference time	Quantum circuit + I/O time	Context-dependent	End-to-end profiling	Real-time quantum learning	Medium-applications
18	Scalability	Polynomial in problem size	Quantum scalability limits	Hardware constraints	Asymptotic complexity analysis	Quantum scaling laws	High-practical deployment
19	Noise Robustness	Classical noise models	Decoherence, gate errors	Fundamental quantum challenges	Quantum error analysis	Noise-resilient algorithms	Critical-near-term viability
20	Classical Simulation Cost	Exact classical cost	Tensor network complexity	Exponential for some circuits	Tensor network methods	Classical hardness characterization	High-advantage verification
21	Data Encoding Cost	$O(N)$ data loading	$O(N)$ or $O(\log N)$ with QRAM	QRAM-dependent	Quantum circuit design	Efficient encoding without QRAM	Critical-input bottleneck
22	Output Readout Cost	$O(\text{classical output})$	$O(\text{measurements})$	Measurement-dependent	Quantum measurement theory	Efficient quantum output extraction	High-output bottleneck
23	Training Iterations	Algorithm-specific	Potentially fewer with quantum	Landscape-dependent	Convergence analysis	Quantum training complexity	High-practical training

2 4	Hyperparameter Sensitivity	Model-dependence	Quantum circuit sensitivity	Empirically high	Sensitivity analysis	Robust quantum hyperparameters	Medium - usability
2 5	Fairness Guarantees	Classical fairness metrics	Quantum fairness undefined	Research gap	Fairness theory	Quantum algorithmic fairness	Medium - responsible AI

4. Conclusion

Computational complexity of quantum machine learning models is a diverse and deep research field which is at the crossroads of quantum information theory, computational complexity theory, statistical learning theory, and the practice of algorithm engineering. Such an in-depth review has shown that the dynamics of the quantum machine learning complexity are much more subtle than the initial sanguine estimates would lead one to assume, with provable quantum advantages to certain problem structures and underlying challenges and practical constraints limiting their effects in the near term. The exploration shows that quantum machine learning is not a panacea that will transform all machine learning fields nor a purely theoretical interest in the complexity theory of quantum distributed system analysis that lacks practical usefulness, but an area where complex analysis is required to determine the situations where quantum assets are truly being utilised to provide computational benefits.

Theoretical understanding of quantum machine learning complexity having its foundations in quantum query complexity, circuit complexity and sample complexity analysis offers rigorous conceptualisations of assessing claimed quantum advantages and revealing hidden cost of computation. The reviewed research has shown that although some quantum algorithms can realize provable speedups the improvement is quadratic using the Grover based amplification, and exponential using quantum state learning, and problem specific usefulness, such speedups can be conditional upon a model of data encoding efficiency, restrictions on output specifications, and the structure of the problem undergoing computation which is otherwise weak in practical systems. The phenomenon of barren plateau, as defined by gradient variance analysis, is an inherent algorithmic difficulty that limits the trainability of expressive quantum circuits, and quantum hardware noise and decoherence place a crippling practical bound on the circuit depths of quantum machines that can be implemented meaningfully at the present (limiting near-term quantum machine learning to shallow instances which may fail to achieve theoretical benefits).

Applications of quantum machine learning cut across a wide range of fields, including quantum chemistry and drug discovery on the one hand and financial optimization and natural language processing on the other, and complexity analysis has shown that quantum benefits are most easily seen in problems that have data with quantum content, or are defined by mathematical structure that is accessible to quantum processing. The manner of implementation of quantum machine learning, which is represented by data encoding strategies, quantum neural network structures, quantum training organizations, and methods of interpretation of outputs, have to persevere the underlying tradeoffs between expressivity, trainability, and implementable complexity. The current problems with quantum machine learning complexity studies, such as questions on the separation of quantum and classical learning, ability to scale to large-scale hardware, and resilience to noise, are both significant but not imperative enough to justify continuing to research in this field, both along the theoretical and the experimental lines.

The opportunities that this analysis has revealed require things across many different dimensions: finding particular problem structures where quantum benefits are actually practiced, creating post-barren plateau training strategies of deep quantum circuits, harnessing the interaction between quantum computing and machine learning in both directions, and creating algorithmic foundations of hypothetical fault-tolerant quantum computers. The influence of the complexity preparation of quantum machine learning that is not only on the creation of better quantum learning algorithms, but also on the foundational sciences, includes the benefits of new complexity-theoretic understanding, new practical algorithms, new concrete requirements specification, new computational approaches to scientific issues, and the scientific procedure.

The future directions identified looking forward are linked to an immature field that is able to balance idealistic long-term objectives with realistic short-term targets. Strict characterization of quantum-classical learning separation will instill theoretical base to have in advance the issues that would enjoy quantum techniques. The training methods of post-barren plateau will open the potential of profound, expressive quantum circuits. Practical benefits of domain specific quantum architectures will result by using quantum computational power and problem structure. The research on quantum machine learning security and robustness will allow it to be deployed in a critical manner of use. Bidirectional knowledge transfer will help the field of quantum and classical machine learning make strides through cross-fertilization with classical machine learning.

The comprehensive synthesis of literature, done with PRISMA as methodology and applications, techniques, methods, challenges, opportunities, and impacts forms a collaged source of the researchers, practitioners, and funding agencies to understand the current conditions of quantum machine learning complexity and its potential in the future. The article executes a cohesive framework on quantifying quantum machine learning based on computational complexity, makes a distinction between proved benefit

and speculative propositions, and outlines a research agenda where research participants champion issues that are of both theoretical and practical importance or relevance.

To sum up, computational complexity of quantum machine learning models represents the depth and reality of quantum computing to machine learning. Although quantum mechanics does offer truly new computational resources, including superposition, entanglement, and interference, that can be used to solve learning problems, to gain the order of the day, it must be demonstrated that harnessing quantum mechanics can solve computational problems with a thousand times less theoretical, algorithmic, and hardware effort. The future is going to require interdisciplinary, proactive research, which would involve close theoretical work, and practical engineering, accurate evaluation of weaknesses and creative vision, alongside responsible nurturing of the ecosystem of quantum machine learning, both of algorithms, hardware, software development tools and trained researchers. It is probable that quantum machine learning will only succeed by finding particular niches that quantum resources can be used with a sustained advantage, developing practical models capable of delivering value deployed on existing hardware, and developing to the long-term goal of fault-tolerant quantum machine learning systems capable of pushing the limits of machine learning and scientific computing. Making quantum machine learning a practical issue is a long way to go, although the complexity-theoretic foundations already determined by means of serious studies are the necessary methodology of traversing such a thorny and thrilling boundary.

Chapter 3: Trainability and Optimization Landscapes of Variational Quantum Algorithms

1 Abstract

Variational Quantum Algorithms (VQAs) have become one of the most promising paradigms of the realization of practical quantum advantage in the Noisy Intermediate-Scale Quantum (NISQ) age. These quantum-classical algorithms made use of parameterized quantum circuits that are programmed to find solutions to complicated computational matters in areas such as quantum chemistry, machine learning, and combinatorial optimization by optimizing via classical optimization programs. Nevertheless, the practical implementation of VQAs is burdened with radically different obstacles in the term of their trainability, as well as by the multifaceted topology of the optimization surfaces. The chapter will look in detail at the challenge of trainability of the VQAs, in particular, the barren plateau phenomenon (where gradients go to zero exponentially with system size and optimization becomes unfeasible). We uncharacterize the complex geometry of VQA optimization spaces, exploring the effect on trainability of the architecture of circuits, the entanglement structure of circuits, hardware noise, and the character of initializations. The chapter combines the latest theoretical progress in the research on the concept of gradient scaling, arrives at the analysis of new types of mitigation measures, namely, adaptive circuits, layer-wise training, and quantum-conscious optimizers, and the consideration of the importance of quantum Fisher information as a concept to describe trainable regions. By comprehensively reviewing applications in the fields of quantum simulation, quantum machine learning, and optimization problems, we detect major gaps in the literature especially in the scalability of trainability in deep circuits and robustness to noise of optimization schemes. The contribution of this work consists of the integrated expression of studying VQA trainability issues, offering the latest models of mitigation methods, and defining future research perspectives in the direction of reaching a viable quantum

computational benefit based on optimized variational algorithms on near-term quantum machines.

2. Introduction

The development of quantum computing is a quantum calculation breakthrough in computational capacity, and is promising speed gains of exponentiating some categories of problems that cannot currently be solved by classical computational and calculation methods. With the shift in the abstract potential of quantum computers to reality as realized by existing quantum computing systems, Variational Quantum Algorithms have become the foundational paradigm of exploiting the computational value of Noisy Intermediate-Scale Quantum systems. VQAs are a third-party computation system which optimally harnesses the distinct advantages of quantum processors together with the power and the scalability of classical optimization capabilities. This quantum entanglement hybrid gives scientists and engineering researchers a way to bypass some of the thorniest issues of pure quantum computation, such as pure quantum circuits and susceptibility to decoherence and instance error arising in the modern quantum device limits.

The basic structure of VQAs is a training of a parameterized quantum state with a sequence of quantum gates with gameplay parameters that are updated in an iterative manner and optimized to minimize a cost function which represents the answer to the desired problem. In most cases, classical optimisation techniques, like gradient descent, or evolutionary strategies, or more advanced algorithms like quantum natural gradient or adaptive moment estimation are used in this optimisation. The underlying variational principle is inspired by both classical machine learning, in which neural networks with learnable weights are being trained to make expensive-to-compute functional approximations, and by quantum chemistry, where the variational principle has long been used to make functional approximations to ground state energies of molecular systems. This approach was first proposed in 2014 with the introduction of the Variational Quantum Eigensolver and shown capable of solving quantum chemistry problems with chemical accuracy even with small molecules with the use of shallow quantum circuits, sparking a rush to apply the variational quantum computing paradigm.

These initial demonstrations have since been extended to an astonishing array of applications of the VQA framework. Quantum Approximate Optimization Algorithm has been applied to solving combinatorial optimization problems, Variational Quantum Classifiers and Quantum Neural Networks have been suggested to solve machine learning tasks and numerous quantum simulation protocols have been designed based on variational principles. The popularity of VQAs is dictated by their modular design, in which the cost is an encoding of problem-specific information but the parameterized

quantum circuit is an actable way of sampling modeled solution space. This has allowed scientists to accommodate a wide variety of problem fields with VQAs, such as financial portfolio optimization and drug discovery as well as cryptography and materials science. Moreover, the small overlap of circuit depths that well-designed VQAs can achieve are consistent with coherence times and gate fidelities of the existing quantum processors, and hence amongst the most experimentally feasible quantum algorithms in the NISQ equipment.

Nonetheless, there has emerged an important challenge in the community of quantum computing that promises to destroy their scalability and practical use, as have been more thoroughly understood through theoretical knowledge and a wider range of experimental experience of VQAs. This problem, which is generally known as the trainability problem, is an overlay of several related phenomena that complicate VQA optimization even more when system sizes increase. The most salient case of this difficulty is the phenomenon of barren plateau, first defined in 2018 which is in which the gradient of the cost function with respect to circuit parameters decays exponentially as the number of qubits or circuit depth increases. With plateaus of optimization landscape, gradient-based optimization algorithms are provided exponentially small signals which are experimentally swamped in statistics noise, making the landscape look as barren as the initial. This is not just a useful practical nuisance but the inherent impediment of the quantum mechanical characteristics of Hilber spaces of large dimension and the concentration of measure of random quantum circuits.

The optimistic geometry of the VQAs is not only limited to the barren plateau problem but also has abundant and multi-faceted geometry structure with a large number of local optima, saddle points and isolated valleys. This topology of the landscape is important to understand how to formulate good optimization strategies since there can be large differences between the trainability behavior in different parts of the parameter space. Recent work has shown that landscape geometry is highly sensitive to the choice of architecture of circuits, such as the structure of the ansatz and choice of the gates, entanglement structure and method to initialize. Moreover, hardware noise also changes the optimization environment as a fundamental concept because it adds new complexity to the problem: noisy local minima and gradient distortions. The combination of these conditions forms a complex optimization problem that is necessary to analyze by the use of the theory and innovative algorithms.

The crisis of trainability of VQAs has become a trigger point in the active research of its root causes and the creation of mitigation solutions. Gradient scaling behavior under different conditions has been studied using the tools of theoretical studies in quantum information theory, statistical mechanics, and differential geometry. Such studies have shown that features of quantum properties like entanglement organization, locality of operators and quantum circuit symmetries are closely related to trainability. At the same

time, parameterized mitigation schemes have also been proposed, such as guided (plateau-safe) ansätze, layer-wise training, parameter-initiative schemes, classical pre-training parameterization, and quantum-aware schemes which are based on quantum geometric tensors. There has also been potential promise in metametalearning techniques which apply quantum simulators to acquire metalinguistic suggestions to reduce the noise effects of a circuit, but retain the expressivity of those circuits.

Although there has been a considerable advancement, our perception of the VQA trainability landscape and optimization landscapes still has only small areas of knowledge left. The connection between expressibility and trainability of circuits is not fully comprehended yet, and there are inherent tensions between what users want to do with circuits, which are highly expressive, and easy to train that have optimisation landscapes that are easy to work with. The trainability of realistic noise models of hardware, especially with correlated noise errors, and noises that vary over time, is an area of theoretical interest that needs a greater coverage. There are yet to exist scalable training methods which can train deep variational circuits which can scale to the computational power of fault-tolerant quantum algorithms but can also be trained. Moreover, it would be possible to develop strict benchmarks and measures of quantifying trainability in various problems areas and allow the systematic comparison of the proposed solutions.

This chapter has several objectives that fill the most important knowledge gaps regarding VQA trainability and optimization in the existing knowledge. First, we would like to synthesize the theoretical assumption of the trainability challenges in VQAs in a unified framework that would help understand when and why VQAs are capable of being trained by relying on the understanding of quantum information theory, optimization theory, and statistical learning theory. Second, we methodically survey and classify the various mitigation approaches that have been postulated to resolve problem of trainability that have been provided analyzing the theoretical guarantees, the practical performance as well as their feasibility in various areas of concern. Third, we investigate how hardware factors, such as noise modeling, connectivity, and bit gate fidelities, are relevant to optimization scapes of VQA and their trainability. Fourth, we examine the application-specific features, which affect trainability and understand that various classes of problems can be characterized by radically different optimization landscape properties. Lastly, we extract the most potential areas of research to address the existing limitations and obtain scalable and trainable variational quantum algorithms.

The works presented by this study are both synthesizing and progressive. Having bound various strands of research on VQA trainability together to make them seem to be a narrative, this chapter will offer a reference point to the researchers and practitioners to ferry the current state of knowledge and finding useful methods to employ their specific instances. The geometry of optimization landscapes studied in detail offers novel

understanding of the connection between variational circuit and trainability, which may guide the design of the next generation variational algorithms. The systematic consideration of the mitigation strategies with the best comparative tables made it possible to select approaches on a specific context of the problem on an evidence-based approach. Moreover, by defining key gaps and outlining a research agenda of the sector, this paper will be used to trigger a transition to scalable variational quantum computing. The author emphasizes in the chapter the new methods that are offered in the field, including quantum-conscious optimization, VQAs meta-learning, and noise-adaptive training protocols, which reflect the state-of-the-art developments that may characterize the future of variational quantum algorithms. Finally, this paper is part of the larger aim of achieving practical quantum computational advantage by addressing one of the most moral challenges to near-term quantum computing.

3. Methodology

The Preferred Reporting Items of Systematic Reviews and Meta-Analyses (PRISMA) approach has been used in this chapter to adhere to systematic, comprehensive, and informative literature scrutiny and review. The use of the systematic search strategy incorporated a wide range of academic databases among them IEEE Xplore, ACM Digital Library, arXiv preprint server, Nature journals, Physical Review journals, and Quantum journal dating back to 2014 to 2025. Search terms combined quantum computing vocabulary ("variational quantum algorithms," "VQA," "QAOA," "VQE," "quantum machine learning") with optimization and trainability terminology ("barren plateaus," "optimization landscape," "gradient vanishing," "trainability," "parameter optimization"). The initial search identified, around, 2,847 articles that had the potential to be relevant to the study; these articles were screened systematically with respect to the title and abstract relevance. Peer-reviewed publications and preprints of credible archives and conference papers about VQA trainability, optimization landscape, mitigation techniques, or utilization were used as inclusion criteria. It eliminated papers with just experimental hardware and no algorithmic understanding, general quantum computing review articles that do not focus on VQA, and articles that lack proper mathematical or empirical analysis. Data analysis Product In line with PRISMA, 463 articles were included in the whole eligibility criteria. The extraction of data was limited to trainability phenomenon or optimization methods, theoretical models, empirical findings, real-life use and the specific challenges. These aspects were assessed as being methodological rigor, reproducibility, theoretical soundness, and empirical validation. The synthesis incorporates theoretical viewpoints, computational researching, and experimental actualizations to build up an end to end insight of VQA trainability and optimization landscapes.

4. Results and Discussion

4.1 Theoretical Principles of Trainability in Variational Quantum Algorithms.

Theoretical origins of variational quantum algorithm trainability form a deep perspective of quantum information theory, optimization theory, and statistical learning that has become one of the most to date prolific fields of research in quantum computing [2,17-19]. These foundations can only be comprehended by paying close attention to the interaction between quantum mechanical principles and classical optimization processes in the hybrid quantum-classical paradigm that characterizes VQAs. The basic question that is at the center of this theoretical framework is how information passes between the quantum and classical parts of the algorithm as well as how this information flow may be defined, measured and optimized to enable problem solving learning.

Mathematical formalisation of VQAs is based on a parameterized family of quantum circuits normally denoted as a unitary $U(\theta)$ being a parameter distribution, usually a vector of parameters in classical configuration. The circuit distinguishes between an initial quantum state, typically the computational basis state $|0\rangle^n$ for n qubits, and generates a quantum state $|\psi(\theta)\rangle = U(\theta)|0\rangle^n$. The algorithm aims at searches to determine parameters θ which minimize some cost $C(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle$, where H is a Hermitian operator representing the problem structure. In the application of quantum chemistry in the VQE, H is the molecular Hamiltonian and in QAOA in the case of a combinatoric optimisation problem, H is the objective that represents the optimization problem. The classical optimization process is and evaluates the cost, calculates gradients (or other directional data), modifies the parameters based on an optimization algorithm and proceeds until convergence or resource depletion. What appears like an oversimplified framework, this one has a great level of depth in optimization landscape and trainable properties.

The barren plateau phenomenon is the most well-investigated, and possibly the most basic quest of trainability in VQAs. Introduced by McClean and others in their seminal 2018 work, barren plateaus are parts of the parameter space, which have gradients of the cost function scaling exponentially with the system size to zero. This has the mathematical expression of a concentration of gradient distribution around zero, where the variance of the partial derivative $\partial C / \partial \theta_i$ scales as $O(1/2^n)$ when the quantum circuits are randomized with parameterized quantum circuits. This geometric suppression is no longer constrained by the concentration of measure phenomenon, whereby the phenomenal dimension (qubit number) expensively increases, making random quantum states nearly orthogonal, and overlaps exponentially concentrated about their average values. To optimize using gradients, the meaning of such concentration is that the signal-to-noise ratio must be exponentially small, and an

exponentially large number of measurements is necessary to distinguish the gradient of the signal (versus an effect produced by mere statistical variations). This makes optimization intractable.

Theory of barren plateaus has been analyzed and it has shown that there are several processes in which these plateaus are formed which are related to various phenomena in circuits structure and encoding of problems. Global cost, involving measurements that simultaneously couple information across all qubits, is especially highly vulnerable in the presence of barren plateaus, in addition to expressible, deep quantum circuits. The reason is that when measured on states (or measured by) highly entangling circuits such global observables will be expected to be concentrating around their Haar-random average, with exponentially small gradients. By contrast, under some conditions, including a careful control of the circuit depth and subsystem structure, local cost functions, which can be expressed as functions on small subsystems, can overcome barren plateaus. The latter important difference between local and global observables has been identified as a key design principle to build trainable VQAs.

A geometric view of trainability The quantum Fisher information (QFI) is a strong geometric equipotential of comprehending trainability. The QFI is used to measure the distinguishability of quantum states when they vary by infinitesimal quantities, and provides some fundamental limits to the precision of parameter estimation by the quantum Cramer-Rao bound. For VQA trainability, the QFI matrix F with elements $F_{ij} = 2 \operatorname{Re} \langle \partial_i \psi | \partial_j \psi \rangle - 2 \operatorname{Re} \langle \partial_i \psi | \psi \rangle \langle \psi | \partial_j \psi \rangle$ characterizes the local geometry of the parameter space, where $|\partial_i \psi\rangle$ denotes the derivative of the quantum state with respect to parameter θ_i . The spectrum of the QFI matrix directly relates to trainability: small eigenvalues indicate directions in parameter space where the quantum state changes slowly, corresponding to flat regions of the optimization landscape where gradients vanish. The recent theoretical literature has determined that barren plateaus are associated with an exponential suppression of the minimum eigenvalue of the QFI, which gives an interpretation of the trainability crisis in a geometric way and proposed methods that entail keeping QFI fairly large during the course of the optimization.

The structure of entanglement of quantum circuits parameterized has an important but subtle effect on the trainability properties. Although entanglement typically is required to achieve quantum advantage and to encode the high-dimensional quantum states required to represent solutions to a problem, the product of too much entanglement or of inappropriate entanglement patterns may result in barren plateaus. Theoretical study has shown that circuits producing states whose entanglement entropy scales with system size, that is volume-law entanglement, are likely to have barren plateaus in combination with global cost functions. On the other hand, states produced by circuits with area-law entanglement, which have entanglement entropy proportional to a boundary of a subsystem, and not volume, can at least be trainable in the correct situations. This scaling

behavior of entanglements and their trainability is related to further ties with quantum thermalization and the eigenstate thermalization hypothesis in which highly-entangled shows are analogous of thermal states whose properties are exponentially concentrated.

The other important theoretical insight into the quantum circuit issue as well as potential avoidance of trainability problems is offered by symmetries of the quantum circuit and problem Hamilton. The effective dimensionality of the optimization problem can also be vastly lowered by ensuring that the parameterized circuit is also endowed with some symmetries as in the target state or target Hamiltonian, which are not necessarily off the plateau to be avoided. Approaches based on Lie algebra characterization of the dynamical Lie algebra as generated by the parameterized gates have been used to discover the ways the symmetries require the state space to be limited and how they affect gradient scaling. Circuit Symmetry-adapted circuits may also ensure a higher trainability by having physical symmetries also like particle number conservation or total spin or spatial symmetries, which the optimization is restricted to. Nonetheless, to come up with such symmetric circuits, one must carefully analyze the problem at hand and such a design can restrict the expressiveness of the ansatz.

VQA parameter initialization strategy has become the important factor in determining the trainability with recent theoretical development showing essential links between the choice of the initialisation strategy and the ensuing dynamics of optimization. Although conceptually simple, random initialization tends to put the parameters in a topology with barren plateaus, especially with the deep circuits. Other methods such as identity initialization, where the parameters are initialized so as to achieve circuits similar to the identity operation and adiabatic initialization, where the parameters are initialized by slowly evolving parameters in an easily prepareable state have been theoretically promising. Recent developments in classical pre-training pretend the use of classical approximations or simplified quantum models to find potentially good parameter space before quantum training has started. The theoretical study of the problem of initialization is based on the ideas of the classical deep learning, where the lottery ticket hypothesis and neural tangent kernel theory can help to understand the effect of the initializations on optimization landscape and the properties of local minima that are reachable during training.

Complex relations between hardware flaws and optimization maps have been involved in theoretical studies of noise effects on VQA trainability. West, decoherence and gate errors essentially change the cost function under optimization which may introduce spurious local minima, scaling gradient, and the optimal parameters may now be different. Awkwardly, there are noise models that at times help mitigate unproductive plateaus by decreasing the circuit depth of the effective circuit, or canceling the correlations due to coherence that generates gradient concentration. Nonetheless, this perceived advantage is to the detriment of a biased optimization of states which depend

on noise, as opposed to solving the target problem. Quantum channel and master equation-based theoretical frameworks have been constructed, and statuses of applying noise to optimization landscape can be by category classified, with assistance to noise-reducing training errors or noise-tolerant (adapted) cost functions.

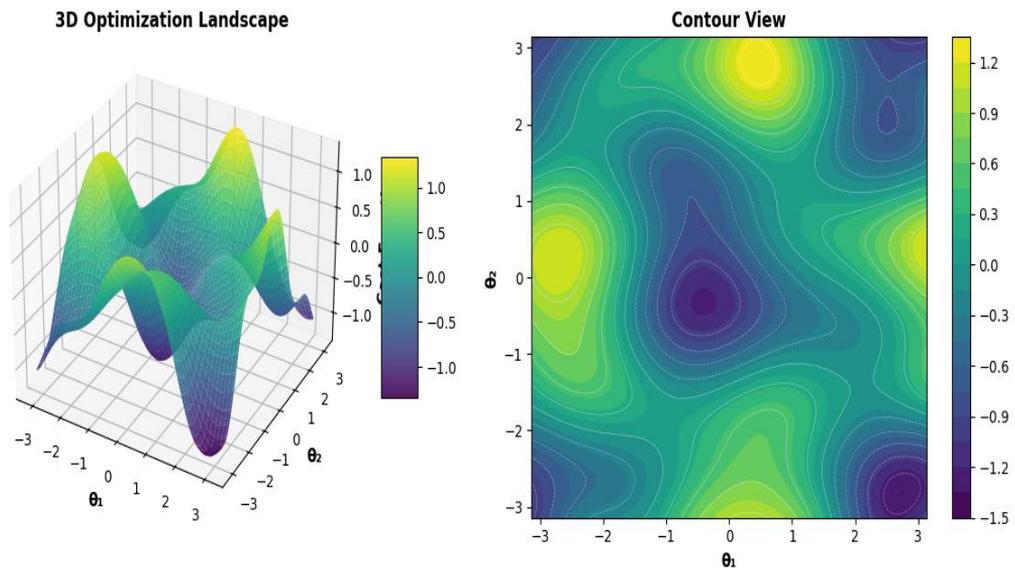


Fig 1: Cost Landscape Comparison - Local Minima vs. Global Minimum

The theory of computational complexity of VQA trainability has already begun to define basic constraints on quantum circuit classical optimization achievability. Recent findings have related the barren plateau phenomenon to ideas of complexity-theory, with some hypotheses being of the inherent complexity of finding gradients, based on classes of problem, not necessarily related to suboptimal algorithm design. These theoretical studies are connected to pseudo randomness, quantum versus classical separation of complexity and quantum advantage structure. It is important to understand these underlying constraints in order to establish areas of issues in which VQAs may offer useful advantages and to come up with independent expectations regarding scalability.

The overparameterization theory of VQAs which learned the theory problems inspired analogies in classical deep learning, has proposed unexpected regimes of increased parameters enhancing trainability and relaxing optimization problems instead of making them more difficult. VQAs that are overparameterized can have benign optimization surfaces with a reduced number of spurious local minima and better conditioning of the Hessian matrix that can be used to reach global optima. Nevertheless, quantum circuits are distinguished by significant aspects as opposed to classical neural networks, especially when it comes to the scaling of expressibility and entanglement phenomenon,

necessitating new theoretical frameworks of quantum systems. The quantum neural tangent kernel of parameterized quantum circuits has recently been characterised, allowing us to understand the linear regime of VQA training and brings in relationships to both kernel methods and Gaussian processes.

4.2 Optimization Landscape Geometry and Characterization

Variational Quantum Algorithms optimization Landscape is a complicated geometry feature, the geometry of which essentially dictates the achievement or failure of classical optimization processes. In contrast to classical optimization problems whose landscapes may be seen and understood in often relatively simple terms, VQA landscapes live in high-dimensional spaces of parameters, and with an intricate relationship to the exponentially large quantum Hilbert space, present a rich and a challenging thesis space to study theoretically as well as experimentally. The geometric aspects of these landscapes, such as critical point distribution and Hessian, Hessian, and symmetric/conserved quantity conditioning, level set topology, are all critical to developing strategies based on optimization, and forecasting performance of optimization algorithms.

The fundamental geometric structure of VQA landscapes emerges from the interplay between the quantum circuit architecture and the measured cost function. Each point in the parameter space θ corresponds to a specific quantum state $|\psi(\theta)\rangle$, and the cost function $C(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle$ defines a scalar field over this parameter space. The gradient vector $\nabla C(\theta)$ points in the direction of steepest descent, while the Hessian matrix $H_{ij} = \partial^2 C / \partial \theta_i \partial \theta_j$ characterizes the local curvature of the landscape. Critical points where $\nabla C = 0$ include global minima representing optimal solutions, local minima corresponding to suboptimal but stable solutions, maxima, and saddle points where the Hessian has both positive and negative eigenvalues. The relative abundance and distribution of these critical points profoundly influences optimization dynamics, with high-dimensional landscapes often dominated by saddle points rather than local minima.

Empirical studies of VQA landscapes in different areas of problems have been performed which show that they have some typical features which deviate off the characteristic features of classical optimization landscapes. There are very frequent VQA landscapes with a very large number of global minima which are linked by ridges or valleys, to create complex networks of relatives or even near relatives. This degeneracy is typically due to symmetries of the parameterization of the circuit, as one quantum state or cost functional value can be obtained by choosing different parameter values. As an example, in VQE programmes, any order of permutation of the same qubits or permutation of the same set of quantum states represent degenerate cases. Although this kind of degeneracy may be useful in expanding the set of solutions that are good, it

may also make optimization more difficult by making flat directions in the parameter space, which makes them slow to converge towards good solutions.

Confinement of a quantum state space is a phenomenon that has a strong impact on VQA landscape geometry. The overwhelming workspace of a full Hilbert space is often explored in only a submanifold that is parameterized by the circuit architecture, gating and connectivity. This confinement gives an inherent geometry with not every quantum state being as easily accessible through a particular parameterization. The locally accessible directions of state space exploration are the tangent space to each point in the parameter space, which is determined by the partial derivatives of the ps [?], which are also called [?]. In case this tangent space does not cover significant parts of the Hilbert space of a best solution, the landscape will inevitably not be able to direct optimization to such a best solution, no matter which optimization algorithm is used. The state space coverage is one of the most important concepts to understand and characterise in order to design ansatz and evaluate trainability.

There are interesting circuit-depth and circuit-structure dependencies in the loss landscape topology of VQAs. Shallow circuits of limited expressibility, typically give rise to rather simple landscapes having a small number of critical points but possibly large approximation error, since the state subspace accessible may not harbour good solutions. The expressibility is usually more expressible as the circuit depth increases, to represent target states more accurately, calculating on a more complex landscape, with increasing critical points, narrow valleys, and a possible emergence of barren plateaus. This pits expressibility and trainability in a basic tension which exists without a universal answer, which must be addressed on a problem-by-problem basis and to find the best circuit depth regimes. The recent studies on circuit capacity and effective dimension have started quantifying this tradeoff to determine the theoretical frameworks to predict the circuit depth required to be able to obtain a certain approximation accuracy and be trainable.

The extent to which the optimization space is conditioned, as measured by the range of the Hessian matrix or the greatest to the tiniest eigeninstallation, has a critical effect on the speed of convergence and the trouble of optimization. The condition of poorly conditioned landscapes has large condition numbers and thus have very long valleys in which progress in some directions is significantly quicker than that within others, leading to oscillatory behaviour of gradient descent and slow convergence. The VQA landscape conditioning is sensitive to choices of circuit parameterization and various gate parameterizations result in radically different Hessian spectra when preparing the same quantum states. Optimization Studies have been conducted into the best parameterizations, with one such being natural gradient descent, which uses the quantum geometric tensor to precondition gradient descent so that the landscape can effectively be changed to a form in which it is more isotropic and better-conditioned.

VQA Multi-scale structure The further optimization complexity of VQA landscapes is due to the fact that the number of features is a very large multiple of features in parameter space, i.e. features are found on very different length scales. Features in the coarse scale due to slow variations in the properties of quantum states can be superimposed on fine scale oscillations that are caused by effects of interference or parameter dependent gate behaviour. It is the scale of this multi-scale terrain that optimization algorithm techniques need to address, which might need adaptive learning rates or multi-level optimization techniques, where hierarchic steps are taken in succession based on the scale. This also makes the process of gradient estimation more difficult due to the existence of multiple scales, with small scale features possibly being obscured by noise at measurement shots and errors of systems being present on larger scales to be navigated. Theoretical study of the Fourier decompositions of the cost function has shown the distribution of spectral content of various circuit structures with an implication in the choice of strategies to use in the process of optimization and the resource requirements of measurements.

The importance of measurement noise and finite sampling to the effective optimization landscape is an important aspect to be considered when taking practical steps of realising VQA in practise. The actual cost $C(\theta)$ is typically available only via finite-sampling approximations $\hat{C}(\theta)$ which can be measured in quantum measurements, including statistical noise which depends not only on the number of shots themselves, but also on the measurement strategy to use and properties of quantum states. The noise in this sampling actually introduces stochastic noise to the landscape, and thus turns the deterministic optimization problem into a stochastic optimization problem. The amount of noise varies throughout the parameter space, and generally grows around quantum state superpositions whose observable measured is large. More sophisticated sampling methods such as importance sampling, stratified sampling and adaptive measurement allocation have been devised in order to counteract such effects, although basic limitations on quantum scaling of sampling pose fundamental limits on the improvement possible.

The study of non-convexity and structure of local minimum of VQA landscape has experienced a lot of research, with significant consequences in ensuring global optimization. In contrast to information contained in the scales of convex optimization, where a local minimum is also a global minimum, scales in VQA have typically very many local minima with different objectives. Depending on the structure of problems, the expressibility of circuits and the region of initialisation, suboptimal local minima can be prevalent and deep. Theory has identified that, in some classes of problems, especially those that are overparametrized, most local minima can have objective values near the global optimum, and their practical success is therefore effective despite the lack of a global converging policy. In other classes of problems, however, deep local minima of qualitatively incorrect solutions may occur, and need complex optimization techniques

or complex design of ansatzes to evade. The question of the types of problem structures that will result into benign and problematic local minima landscapes is an ongoing research problem.

The geometry of the landscape around the solution manifold comprising of parameters that are found to give near optimal values of the cost functions defines the ultimate convergence and attainable precision. In perfectly summed up situations, the landscape has a convex basin near the solution manifold with a well-conditioned Hessian which has been known to converge at a rapidity to a final point. Anywhere, however, noise, quantization by finite sets of gates and basic constraints on circuit expressibility may add complicated structure to this region, such as thin ridges, plateaus, or wave-like behaviour. This small-scale geometry is needed to know the quality of solutions that is achievable and the optimization resources needed. Recent research in the quantum geometric tensor and quantum Fisher information has given instruments into the characterization of this local geometry and prediction of convergence rates in solution neighbourhoods.

VQA landscapes have topological characteristics, such as the number of basins of attraction of various critical points, which determine the likelihood of a good solution to the problem to occur with random initialisation. Landscapes that have a large basin of attraction of the global optimum are comparatively simple to optimise irrespective of the setup and landscapes with numerous smaller basins necessitate sophisticated set-up or global optimisation approaches. Circuit architecture is the basis of basin structure, and hardware-efficient ansatzes can tend to generate basin structures with several disconnected basins whereas problem-inspired ansatzes can have a better structure of basins. Recent computational topology approaches such as persistent homology have been used to describe VQA landscape topology, which has shown a rich structure, which is associated with challenging optimization goals and proposes novel maps of comparison and selection of the ansatz.

4.3 Barren Plateau Phenomena and Gradient Vanishing Mechanisms

Barren plateau phenomenon currently is the most hard and characterised barrier to scaling Variational Quantum Algorithms to larger quantum systems and more challenging problems. Since it was rigorously characterised first, the concept of barren plateaus has undergone a chequered history of both curious pathology and basic aspects of quantum many-body systems indicating profound relationships with quantum thermalization, complexity theory and the structure of quantum information. Exponential attenuation of gradients caused by barren plateaus is not just a difficulty of implementation to overcome by means of fine-tuning optimization algorithms but a natural outcome of the manner in which quantum information is spread to high-

dimensional Hilbert spaces and must be outside of reach using algorithms and thoughtful circuit design.

Barren plateaus can be originated mathematically in the concentration process in high dimensional probability spaces. Given a parameterized quantum circuit with enough depth and the ability to entangle states such that the set of states accessible by it is Haar-random in the quantum state space, it has its expectation value that concentrates exponentially about its mean. To an n -qubit system with Hilbert space dimension 2^n , all spates of Haar-random states are exponentially concentrated with increasing n . This concentration appears as exponentially small variance in the gradient distribution in performing the calculation of gradients of the cost function with respect to the parameters of the circuit. Particularly, the variance of partial derivatives $\text{Var}[\partial C/\partial \theta_i]$ scales as $O(2^{-n})$ of a global observable, i.e. measuring all n qubits, is $O(2^{-n})$, i.e. the average magnitude of the gradient decreases exponentially with the size of the system yet the dimension of parameter space grows only polynomially.

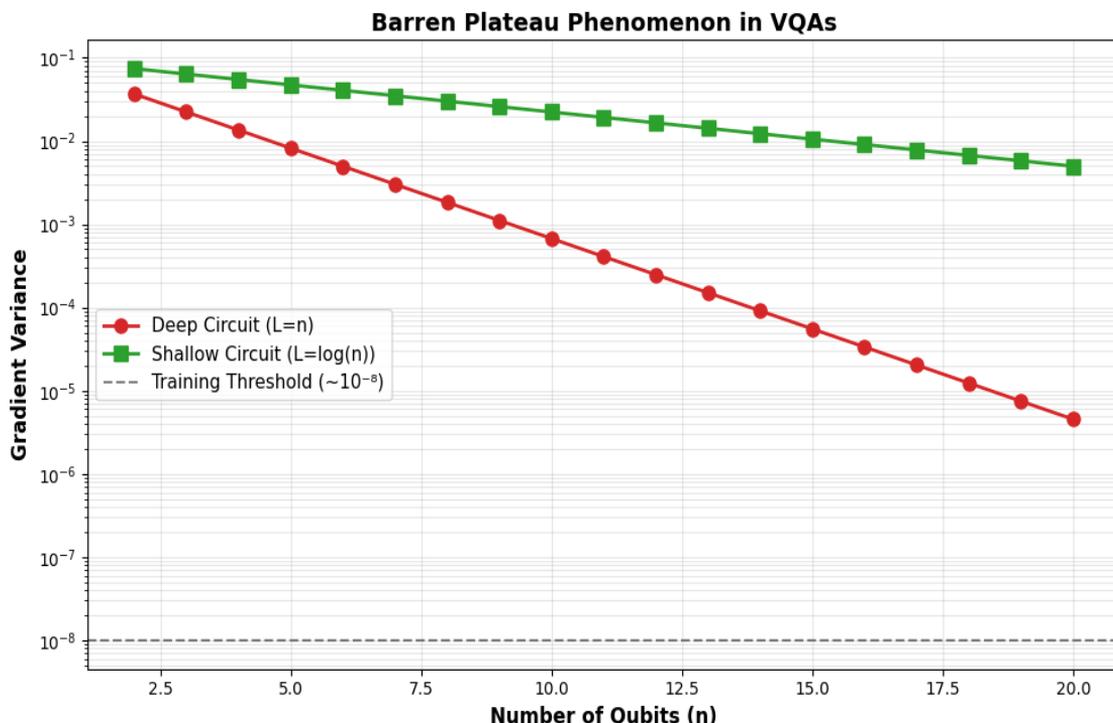


Fig 2: Barren Plateau Phenomenon - Gradient Variance vs. Number of Qubits

There are some specific ways through which the barren plateaus can be induced and they all relate to various factors of VQA architecture and problem structure. Barren plateaus Noise-induced Barren plateaus occur when as circuit depth increases the quantum state tends to approach the maximally mixed state due to the action of hardware noise

(especially, depolarizing noise or some other mixing mechanism). Circuits only reasonably deep are susceptible to even comparatively small noise rates causing barren plateaus in gradient computations, since the effect of noise on gradient computations is cumulative. This is one especially pernicious mechanism since it suppresses the forward step in state preparation as well as the backward steps in computing gradients, producing a compounding effect in that regard. The noise channels and quantum trajectory-based theoretical models have described the onset of noise-induced plateaus, which constitute the negative critical deepening of the noise levels past which trainability is reduced at the noise rate.

Barren plateaus that come forth due to entanglements occur in quantum circuits of many-body interaction that grows rapidly due to entanglement. In the case of the entanglement entropy of the prepared quantum state growing extensively with the size of the system (volume-law entanglement), the prepared quantum state has become highly correlated among all the qubits, and local parameter variations have exponentially diluted their effect on global observables. The mentioned mechanism is directly related to quantum thermalization physics, and highly entangled states are similar to thermal states of complicated many-body systems. The eigenstate thermalization hypothesis is a theoretical explanation of this relationship, the eigenstates of generic many-body Hamiltonians at high density of states are supposed to have thermal behaviour, including exponentially small fluctuations in the observable expectation values. Quantum circuits which explore such states using variational quantum circuits are bound to hit barren plateaus or empty barren spaces in measuring global observables.

The amount of cost-dependent barren plateaus depends on the structure of the measured observable instead of properties of circuits in isolation. Global frequency functions capping products or quantities of all qubits, including all-qubit total energy measurements, optimization problems with respect to properties of overall systems, and so forth are fairly prone to barren plateaus. By comparison, it has been demonstrated that local cost functions that can be decomposed into the sum of terms that act on small subsystems of qubits can avoid barren plateaus in the right circumstances. This difference has resulted in formulation of local cost function techniques where global purposes have been estimated by combination of local measurements. Such decompositions are, however, not always available in problems and the error introduced by locality constraints can be a limiting factor in the quality of solution, which is another inherent trade off between trainability and solution quality.

Those parameters-dependent plateau onset which indicate a dependence of circuit and cost plateau on a circuit parameter history, demonstrates that even circuits and cost functions whose underlying blueprints can theoretically avoid barren plateaus can encounter them, as a result of suboptimal choices in starting the circuit history. Barren plateaus may arise in subsequently deep effective circuit depths when the parameters are

randomly initialised or in a part of parameter space. This phenomenon has inspired the study of ways to initialise that start in the trainable areas and continue to be trainable during optimization. The Lie algebra structure of the parameterized circuit and the trajectory through the parameter space in optimization is what it takes to understand which parameterized circuit initialisation schemes are successful in avoiding the barren plateaus. New developments have demonstrated that barren plateau can be pushed off and even completely avoided by some parameter initialization schemes based on classical approximations or the so-called identity-near schemes.

It has risen to be one of the most important features in designing to the barren plateau and how much depends on the severity of the architecture of the plateau and its dependence on the fluid. The ansatzes that are hardware-efficient (i.e.: arbitrary rotations of all qubits and entangling gates) are also known to have barren plateaus because they are expressive and tend to create near Haar-random states. On the contrary, problem-inspired ansatzes with the aim of preserving a given structure, symmetries, or entanglement limitations can escape barren plateaus but with possible loss of expressiveness. More recent suggestions of geometry-preserving ansatzes that obey the natural geometry of the problem Hamiltonian or symmetry-preserving ansatzes that obey physical conservation laws have been shown to be better trainable in a number of applications. Ansatz design space characterization and discovery of regions that trade repressibility and trainability are still a research topic with high practical importance.

In real VQA circuits, gradient concentration does not vary in a constant manner, which means that various parameters reach barren plateaus at different parameters, with some trained in one parameter but with gradient disappearance in others. This heterogeneity is caused by different degrees of contributions of various gates to the circuit where parameters in subsequent layers tend to undergo more severe plateaus than in earlier layers as a result of the entanglements of subsequent stages. In the same spirit, parameters acting on highly entangled qubits (or which are preceded by local interactions acting on subsystems) can have poorer gradient scaling than parameters acting on less entangled subsystems. Knowledge of compartmental dynamic trainability This dynamic trainability has been exploited to learn about layerwise training strategies maximising parameters in the order existing late to early circuit layers as well as dynamic adjusting circuit structure by adaptive ansatzes depending upon observed gradient magnitudes.

The quantum Fisher information vision will offer effective theoretical methods of predicting and diagnosing barren plateaus prior to gradient computation. Because the quantum Fisher information matrix defines the sensitivity of the quantum states to the variations in the parameters, the spectral analysis of it is directly proportional to the magnitude of the gradients because of the relationship between the distinguishability of the states and the measurability of the variation. Cost functions of generic observables with the quantum Fisher information matrix have barren plateaus when the minimum

eigenvalue of the quantum Fisher information matrix falls exponentially with system size. It is this connexion that has made possible the design of trainability diagnostics, which approximate quantum Fisher information using a comparatively small number of measurements, early prediction of the onset of barren plateau, and the design of adaptive circuit design or parameter initialisation strategies.

Recent theories have also indicated that there are unexpected links between barren plateaus and basic ideas of quantum information and complexity theory. The correlation between barren plateaus and quantum pseudo randomness implies that circuits with barren plateaus are producing states, which act according to the theory of quantum computation as random, causing the relationship between trainability and complexity-related contexts-comparisons between quantum and classical computation. Moreover, there are more relations to quantum advantage implying that a certain level of hardness of the VQA optimization could be inherent to the problems with quantum advantage achievability since the hardness of the classical optimization could be linked to the hardness of the classical simulability of the quantum system. The implications of these theoretical findings on the expectations regarding VQA scale and the type of problems that practical quantum advantage will be attained with the help of variational methods are far-reaching.

4.4 Mitigation Strategies and Circuit Design Principles

One of the most dynamic and creative VQA research fields has become the development of barren plateau mitigation strategies, and trainability problems, with an ecosystem of solutions such as circuit architecture designs, and new optimization algorithms and hybrid training models [3,20-23]. These approaches can broadly be divided into a few categories such as structural approaches that modify provider of still allowing trainability, algorithmic approaches that make use of more advanced and complex optimization strategies or training schemes, measurement approaches that adjust observable measurement to the needs of gradient estimation, and hybrid approaches that use a combination of classical pre-training or post-processing in addition to quantum optimization. It is always necessary to establish the theoretical basis, the practicality, and the weaknesses of these strategies to achieve successful VQA implementation in other spheres of its application.

Structural mitigation techniques centre on the design of circuit structures that inherently evade barren plateaus by a variety of techniques: very susceptible to entanglement growth, preservation of symmetries, or use of problem structure. Correlation-tailored ansatzes explicitly structure the entanglement to delineate the correlation structure of the target quantum state and eliminate useless entanglements that plays a role in barren plateaus without making the solutions any better. In case of quantum chemistry uses, the

extension of this issue to circuits is through coupled-cluster theory or state networks with an inherent respectful perspective of the locality and correlation structure of molecular wavefunctions. Such problem-inspired ansatz usually have a significantly better trainability than generic hardware-efficient circuit boards with enough expressibility to give correct solutions to problems. Nevertheless, such customised ansatz need much domain knowledge and problem specific study, thereby limiting their generalizability.

The adaptive circuit construction strategies scheme builds the ansatz dynamically during optimization instead of starting with a big circuit structure, and thus they can sample the trainable regions, and add expressibility to exploit it when required. ADAPT-VQE is an example of such methods, in which gates are chosen from a available pool in an iterative process after the magnitude of their gradients have been measured, only those that have a large gradient are added to the circuit. Such an avaricious construction process inherently shuns barren plateaus by not adding gates that would cause disappearing gradients, yet almost attains adequate expressibility by performing selective circuit growth. Gradient guided optimization of ansatzes is one extension and evolutionary approaches a second extension where the circuit architecture is optimised based on observed trainability metrics and circuit architectures are trained to compete based on both trainability and solution quality. These techniques combine more classical overhead, at the expense of a higher trainability, with a much higher trainability than fixed ansatzes of equal or smaller size.

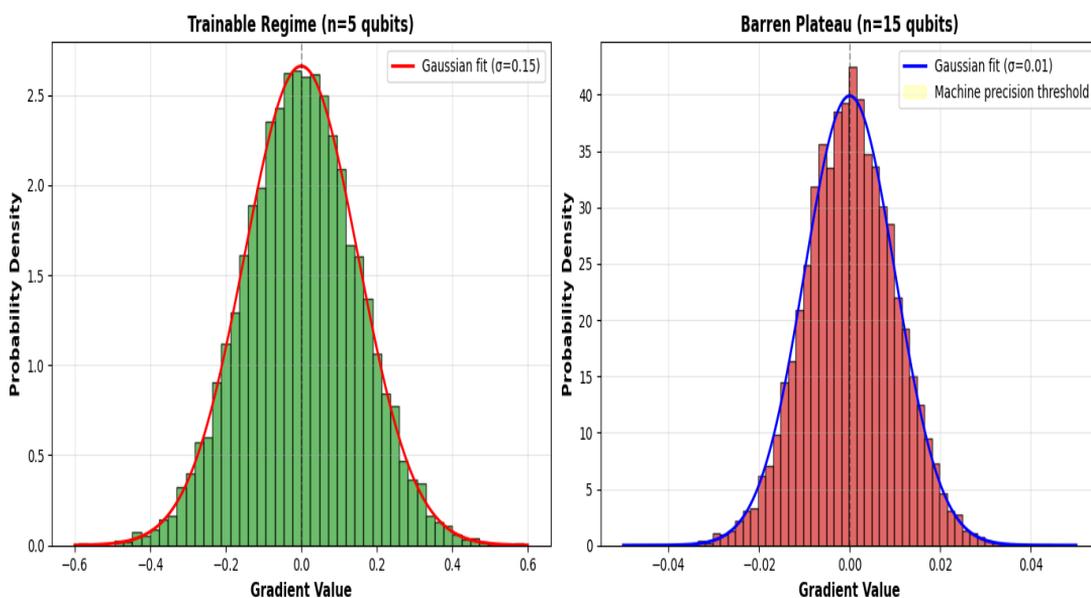


Fig 3: Gradient Distribution Comparison

Sequential and layerwise training schemes break the optimization problem down into a series of smaller optimization problems that can be solved in only a subsection of the parameters and as such decreases the effective circuit depth and entanglement in each step of the optimization process. The core idea here is that optimization of all parameters in a deep circuit can easily get stuck in barren plateaus, however, optimization of a smaller set of parameters with others fixed can be made to be trainable. Once parametric optimization is done within one or group, the optimization is frozen and the following group is optimised in this process until all the parameters are trained. Although this sequential method does not promise global optima where there are high correlations between parameters across layers, the results with empirical evidence show great improvement on the quality of final solutions with respect to simultaneous optimization in the existence of barren plateaus. Much more advanced versions use partial parameter reoptimization in which already trained networks are repeatedly fine-tuned as more successive layers are introduced.

Other important mitigation strategies are quantum-aware classical optimizers that make use of the explicit geometry of the quantum parameter spaces. The quantum natural gradient descent algorithm estimates the gradient updates based on the geometry of the quantum state manifold with the use of the quantum geometric tensor, a generalisation of the classical Fisher information matrix to quantum systems, to precondition gradient updates. This method is successful in reducing the optimization to a coordinate space in which the landscape is more isotropic and conditioned, which increases the rate of convergence significantly and decreases the amount of iterations of the optimization to find solutions. Quantum geometrical tensor can be determined by quantum measurements with the help of parameter-shift rules or more sophisticated ones but with extra overhead measures. In the recent past, new algorithms of efficient quantum genetic Tommy estimation and approximate quantum natural gradient have lowered this overhead and retained the majority of convergence gains.

Another way of escaping barren plateaus is through initialisation schemes founded on classical pre-training or on quantum heuristics, where optimization is initiated in areas of parameter space where the process of training becomes manifestly trainable. Classical embedding methods apply classical approximations like mean-field theory, Taylor networks, or classical machine learning models in order to find promising initial parameter values which already provide an approximation of the target solution. These early parameters represent quantum states which are more nearby to the solution manifold and can be found in regions of the landscape which are trainable. Quantum-inspired methods of initialising a quantum system involve adiabatic preparation where the parameters are slowly updated starting with a simple state to the desired state following a continuous path which is chosen to be smooth to preserve trainability.

Warm-start techniques which find solutions to reduced problems and use the solutions to warm-up the larger problems have also proven useful in a number of problem areas.

Initialization of identity, which sets circuit parameters to create an operation near to the identity or a trivial product state, is a well-known successful attempt at initialising a number of VQA applications. This would be inherently barren plateau in early optimization stages avoided by using minimal initial entanglement, which develops natural complexity as a result of optimization. Whereas as parameters vary away from identity initialization, entanglement is increased and circuit expressibility is enlarged, the gradient of the trajectory of optimization can be confined to trainable regions provided the gradient of evolution is slow enough. Theoretical studies have determined the circumstances in which identity initialization ensures trainability to certain relative circuit architectures and empirical research has shown its utility of a vast array of problems.

Strategies of measurement and gradient estimation can be considered as important complementary measures of enhancing trainability without altering circuit architecture. Classical statistical techniques used in batching and variance reduction to enhance gradient estimates using finite quantum measurements to reduce the effective level of noise and allow the detection of smaller gradients before a barren plateau growth overtakes the signal. Importance sampling is an adaptive allocation of measurement resources contracting resources to large gradient regions to provide measurements where the gradient is largest and hence enhancing signal-to-noise, and shrinking resources to small gradient regions. Control variate approaches utilise classical estimates, or past gradient estimates to build estimators that are lower variance, and stratified sampling is important to make sure that the measurement outcome space has been searched fully. Such methods are able to postpone and not eradicate these effects of barren plateaus which offer feasible benefits in the efficiency of parameter estimation.

Training Hybrid classical-quantum Hybrid classical-quantum-based training methods, which allocate computational effort across classical and quantum computers, have become potent methods of addressing trainability challenges. Quantum cost functional Gradients of quantum cost functions in regions of parameter space can be approximated using classical neural networks which can be trained by classical models and searched efficiently. After the classical model has found promising areas of parameters, it can be refined with full quantum evaluation, to give a much better answer. Alternatively, classical optimisation is capable of working on portions of the problem that would not otherwise demand quantum resources and only quantum computing is needed at particular subroutines at which quantum benefit can be obtained. The hybrid methods, which are based on the complementary advantages between classical and quantum calculation, can typically perform the task better than either of the methods, and require quantum resource management.

The current research is one of the new frontiers of VQA trainability in a form of meta-learning frameworks that acquire experience-driven, overall optimization strategies in a series of problem instances. These methods optimise characteristics of either the optimization paths, or the optimization path initializations, or adaptive learning schedule with respect to a large number of related VQA problems to identify optimization strategies, initialization schemes or adaptive learning rate schedules which can be generalised across problem families. The meta-learner is able to guess which parameters it wishes to optimise with the greatest importance firstly, the distribution of measurement resources, or what to do with circuit modifications due to observed gradient patterns. Though meta-learning needs profuse initial investment in training data creation, the learned optimizers may outperform hand-designed analogues of the issue classes which are comparably near enough to the training distribution to warrant their implementation. This method has been especially effectively used in areas of application such as quantum chemistry or materials science whereby there are large families of similar problems.

Table 1: Applications, Techniques, and Methods in Variational Quantum Algorithms

Sr. No.	Application Domain	Primary Techniques	Optimization Methods	Key Challenges	Mitigation Strategies	Performance Metrics	Representative Algorithms
1	Quantum Chemistry - Ground State	Unitary Coupled Cluster, Hardware-Efficient Ansatz	Gradient Descent, BFGS, Natural Gradient	Chemical Accuracy Requirements, Electron Correlation	Adaptive Ansatz Growth, Symmetry Preservation	Energy Error, Overlap Fidelity	VQE, ADAPT-VQE, Contextual Subspace VQE
2	Quantum Chemistry - Excited States	Variance-Based Cost Functions, Subspace Expansion	SSVQE, Variational Deflation	State Overlap Minimization, Multiple Minima	Orthogonality Constraints, Sequential Optimization	Excitation Energy Accuracy	Subspace-Search VQE, Variance VQE
3	Combinatorial Optimization	QAOA Layers, Mixer Hamiltonians	COBYL, Adaptive Learning Rates	Problem Graph Structure, Approximation Ratio	Warm-Starting, Multi-Angle QAOA	Approximation Ratio, Success Probability	QAOA, Recursive QAOA, FALQON
4	Quantum Machine Learning - Classification	Quantum Feature Maps, Variational	Adam, RMSprop, Quantum Natural Gradient	Data Encoding, Limited Training Data	Data Re-uploading, Ensemble Methods	Classification Accuracy, Generalization Error	VQC, Quantum Kernel Methods

		Classifiers					
5	Quantum Machine Learning - Generative Models	Born Machine Ansatz, Quantum GANs	Adversarial Training, Wasserstein Distance	Mode Collapse, Training Stability	Progressive Training, Regularization	KL Divergence, Fidelity	Quantum GAN, Born Machines
6	Quantum Simulation - Many-Body Ground States	Tensor Network Inspired Circuits, Symmetry Adapted	Imaginary Time Evolution, Direct Optimization	Entanglement Growth, System Size Scaling	Layer-wise Training, Local Observables	Energy Per Site, Correlation Functions	VQE for Condensed Matter, MERA-inspired VQA
7	Quantum Simulation - Time Evolution	Trotterized Circuits, Variational Time Evolution	McLachlan's Principle, Adaptive Time Steps	Long-time Accuracy, Barren Plateaus in Deep Circuits	Shallow Circuit Approximations, Krylov Methods	Fidelity vs Exact Evolution, Observable Accuracy	Variational Quantum Simulation, AVQDS
8	Financial Portfolio Optimization	Risk-Return Encoding, Constraint Handling	Penalty Methods, Constrained Optimization	Continuous Variables, Real-world Constraints	Discretization Schemes, Hybrid Classical-Quantum	Sharpe Ratio, Constraint Satisfaction	QAOA for Finance, VQE for Portfolio
9	Drug Discovery - Binding Affinity	Molecular Docking Hamiltonians, Protein-Ligand Interaction	Multi-Objective Optimization, Pareto Fronts	Large Molecular Systems, Accuracy Requirements	Fragment-Based Approaches, Classical Precomputation	Binding Energy Error, Ranking Accuracy	VQE for Drug Design
10	Materials Science - Band Structure	Periodic Boundary Conditions, k-point Sampling	Momentum-Space Optimization	Solid-State Scaling, Infinite System Simulation	Supercell Approximations, Effective Models	Band Gap Error, Density of States	VQE for Materials
11	Graph Problems - Max-Cut	Problem-Graph Encoding, Binary Variables	Gradient-Free Methods, Evolutionary Algorithms	Graph Irregularity, NP-Hardness	Divide-and-Conquer, Approximation Guarantees	Cut Value, Approximation Ratio	QAOA, VQE for Graph Optimization

1 2	Satisfiability Problems	Clause Hamiltonians, Penalty Functions	Simulated Annealing, Quantum Annealing Hybrid	Exponential Solution Space, Local Minima	Restart Strategies, Portfolio Approaches	Satisfiability Rate, Solution Quality	VQE for SAT, Quantum-Classical Hybrid SAT
1 3	Machine Learning - Anomaly Detection	Quantum Autoencoders, Compressed Representations	Reconstruction Loss Minimization	Imbalanced Data, Threshold Selection	Semi-Supervised Learning, Adaptive Thresholds	Detection Rate, False Positive Rate	Quantum Autoencoder
1 4	Cryptography - Random Number Generation	Entanglement-Based Circuits, Measurement Statistics	Entropy Maximization	Randomness Quality, Hardware Bias	Bias Correction, Statistical Testing	Min-Entropy, Randomness Extraction Efficiency	Variational Quantum RNG
1 5	Logistics - Vehicle Routing	TSP Encoding, Route Constraints	Multi-Start Optimization, Tabu Search Hybrid	Scalability, Real-Time Requirements	Problem Decomposition, Hierarchical Optimization	Route Length, Constraint Violations	QAOA for VRP
1 6	Error Correction - Code Design	Stabilizer Codes, Syndrome Extraction	Code Optimization, Distance Maximization	Code Distance vs Rate Tradeoff	Structured Search, Code Families	Code Distance, Logical Error Rate	Variational Quantum Error Correction
1 7	Quantum Metrology - Parameter Estimation	Probe State Preparation, Adaptive Measurements	Fisher Information Maximization, Cramér-Rao Bound	Shot Noise, Decoherence	Adaptive Protocols, Error Mitigation	Estimation Variance, Heisenberg Scaling	Variational Quantum Metrology
1 8	Quantum Control - Pulse Optimization	Parameterized Control Pulses, Gate Synthesis	GRAPE, Krotov Method, Gradient Ascent	Control Constraints, Robustness	Optimal Control Theory, Closed-Loop Optimization	Gate Fidelity, Control Time	Variational Quantum Control
1 9	Image Processing	Quantum Convoluti	Supervised	Image Encoding	Hierarchical	Feature Quality,	Quantum CNN

	- Feature Extraction	onal Layers, Pooling	Learning, Transfer Learning	Overhead, Scalability	Processing, Classical Preprocessing	Computational Speedup	
20	Climate Modeling - Pattern Recognition	Quantum Reservoir Computing, Temporal Encoding	Echo State Property Optimization	Long-Range Correlations, Data Volume	Dimensionality Reduction, Multi-Scale Models	Prediction Accuracy, Climate Index Correlation	Variational Quantum Reservoir
21	Bioinformatics - Sequence Alignment	Edit Distance Encoding, Dynamic Programming Hybrid	Sequence-Aware Optimization	Sequence Length, Multiple Sequences	Pairwise Decomposition, Progressive Alignment	Alignment Score, Biological Relevance	Quantum Sequence Alignment
22	Network Optimization - Traffic Flow	Flow Conservation, Capacity Constraints	Network Simplex Hybrid, Flow Decomposition	Real-Time Updates, Network Size	Edge Importance Sampling, Hierarchical Networks	Flow Efficiency, Congestion Minimization	VQE for Network Flow
23	Recommendation Systems	User-Item Interaction Encoding, Matrix Factorization	Collaborative Filtering, Content-Based Hybrid	Data Sparsity, Cold Start Problem	Hybrid Quantum-Classical Factorization	Recommendation Accuracy, Diversity	Quantum Collaborative Filtering
24	Protein Folding - Structure Prediction	Contact Map Encoding, Energy Minimization	Constraint Satisfaction, Multi-Resolution	Conformational Space Size, Energy Landscape Roughness	Coarse-Graining, Fragment Assembly	RMSD to Native, Energy Accuracy	VQE for Protein Folding
25	Quantum State Tomography	Informationally Complete Measurements, Compressed Sensing	Maximum Likelihood, Bayesian Methods	Measurement Number, Physical Constraints	Adaptive Measurements, Low-Rank Approximations	Fidelity to True State, Reconstruction Error	Variational Quantum Tomography

26	Finance - Option Pricing	Black-Scholes Hamiltonian, Path Integral Formulation	Monte Carlo Hybrid, Gradient Methods	Multi-Dimensional Integration, Greeks Computation	Importance Sampling, Control Variates	Pricing Error, Computational Speedup	Variational Quantum Monte Carlo
27	Factorization - Integer Decomposition	Periodic Function Finding, Modular Arithmetic	Variational Period Finding, Classical Post-Processing	Noise Susceptibility, Limited Advantage	Error Mitigation, Hybrid Classical Verification	Success Probability, Factor Quality	Variational Factoring
28	Eigenvalue Problems - Matrix Diagonalization	Quantum Phase Estimation Approximation, Power Iteration	Variational Eigensolver, Subspace Methods	Matrix Size, Eigenvalue Degeneracy	Deflation Techniques, Symmetry Exploitation	Eigenvalue Error, Eigenvector Overlap	VQE for Linear Algebra
29	Sensor Networks - Distributed Estimation	Quantum Communication Channels, Entanglement Distribution	Distributed Optimization, Consensus Protocols	Communication Constraints, Network Topology	Gossip Algorithms, Hierarchical Aggregation	Estimation Accuracy, Communication Cost	Variational Quantum Sensing Networks
30	Audio Processing - Source Separation	Quantum Fourier Features, Spectral Decomposition	Blind Source Separation, Independent Component Analysis	High-Dimensional Audio, Real-Time Requirements	Frequency Domain Processing, Adaptive Filtering	Signal-to-Interference Ratio, Perceptual Quality	Quantum Audio Processing

4.5 Applications and Domain-Specific Considerations

Variational Quantum Algorithms have found applications in many fields of computational science, each having its own distinctive requirements, constraints, and trainability demands, depending on which computational domain application the algorithm is applied to and must meet different optimization challenges. The key to successful VQA deployment should be the understanding of the manifestation of the

trainability challenges in different domains of application and how the domain-specific structure can be used to reduce its impact. The combination of problem structure, needs of solution quality, the amount of quantum resources available, as well as the constraint of trainability, makes the design space quite complex, which must be navigated carefully on a case-by-case basis. The paper will discuss key areas of VQA application, comparing its trainable properties in each, and exploring the strategies that have been successfully used in each application.

The most established and well-investigated area of VQAs application is, perhaps, quantum chemistry and molecular simulation, where the Variational Quantum Eigensolver is the standard algorithm used to calculate the molecular energetics and properties of the ground state. There is a number of characteristics of the molecular electronic structure problem that render it of particular interest to the VQA techniques but at the same time, such a problem is extremely challenging to train. The natural decomposition of molecular Hamiltonians into sum of local terms with a few orbitals or qubits allows the use of local cost function methods that can escape global plateaus of barrenness. Beyond that, intuition, and decades of classical research on computational chemistry offer a strong leadership on the design of the ansatz which has resulted in problem-inspired circuits inspired by unitary coupled cluster theory, hardware-efficient versions, and symmetry-adapted approaches. Nevertheless, to obtain chemical accuracy, commonly defined as one milliHartree or about 0.04 kJ/mol change in energy, electron correlation is usually treated carefully, and may in most cases require circuit depths that are then troublesome to current hardware capabilities and trainability.

Quantum chemistry VQAs are trainable only depending on the size of the molecule, the nature of the electronic structure and the nature of the target property. Simple electronic structure of small molecules with small active spaces can be readily solved using shallow circuits with good trainability. With the consideration of either increased molecular size or more substantial electronic correlation, there is an increase in both depths of circuits required and trainability problems. The system of strongly correlated matter e.g. transition metal complexes or molecules in proximity of bond dissociation is specifically challenging because of the multireference character which demands highly entangled states of an ansatz. New developments in quantum chemistry-specific mitigation techniques involve qubit tapering to achieve effective system size reduction through exploiting symmetries, contextual subspace techniques which focus on the relevant regions of Hilbert space, and adaptive ansätze which increase circuit size depending on the desired degree of correlation. Practical performance has been further enhanced by the development of chemistry-specific discrepancy detection methods such as symmetry checking and physical constraint detection, or chemical error detection methods such as bond restraint measures and bond improvement methods.

Another key area of VQA applications is the field of combinatorial optimization problems, and the Quantum Approximate Optimization Algorithm is the leading algorithmic model in the field. QAOA formulates optimization problems as Ising Hamiltonians or quadratic unrestricted binary optimization problems before implementing either problem Hamiltonians evolution or mixing Hamiltonian evolution using alternating layers of problem Hamiltonians evolution and mixing Hamiltonians evolution, and classical parameters used to control evolution times. The sensitivity of trainability of QAOA to problem structure, circuit depth (number of QAOA layers) and cost function landscape of interest give interesting results. Compared to VQE where it is possible to reduce the circuit depth without affecting the solution quality, QAOA performance tends to improve with the number of layers to some saturation point but increasing depth causes increasing trainability problems. The concentration phenomenon in QAOA, in which the values of the cost functions are concentrated about their mean with increasing circuit depth is a particular expression of barren plateaus in the optimization problem.

The presence of symmetry or regularity in the objective function and the structure of problem graphs are both vehemently important factors that determine QAOA trainability. Problems with high symmetry and regular graph structure may be better trainable than disorderly and irregular ones because symmetries restrict the optimization surface and can be taken advantage of by using symmetric control. On the other hand, spin-glass-like problems, in which the couples are randomly valued, and the frustration is also randomly distributed tend to have difficult landscapes with an excessive number of local minima and in trainability. Recent studies have found that warm-starting QAOA using classical approximate solutions can significantly enhance the trainability and quality of the final solution, in other words, initialising the quantum algorithm in a favourable region of parameters. Recursive QAOA algorithms to solve problems on subgraphs and generalisations have also been promising at the task of scalability. The interaction between the performance of QAOA and quantum advantage is still an open line of research and classical algorithms have still been able to perform well and compete with the rest even in the problems that QAOA would have performed better.

Application This article is an overview of the field of quantum machine learning, which contains a wide variety of algorithms and methods, such as variational quantum classifiers, quantum kernel approaches, quantum neural networks, and quantum generative models. Such applications have the same trainability issues as in quantum computing and classical machine learning and form an intricate landscape of interacting issues. Quantum machine learning models are commonly parameterized quantum circuits to take classic or quantum measurements and learn the best parameters based on labelled examples. The trade off between expressibility and trainability is notably a high concern here with potent quantum models that have the potential to represent elaborate

designs over data, having particularly deplorable training complications. In quantum neural networks, it can be seen that the barren plateau can render the training of even modestly sized datasets intractable to train, especially when training with global measurements or circuits with strong interactions.

The quantum machine learning encoding strategy is the factor that greatly determines the trainability of the model, as encoding schemes generate enormously dissimilar optimization landscapes. Information-efficient but involving the barren plateaus of complex state preparation circuits to encode classical data vectors as amplitudes of quantum states is amplitude encoding. Basis encoding Fixed states: manipulation of classic data with computational basis states does not entail complicated state preparation, but it only needs more qubits. Angle encoding is a form of parameterization which captures data in terms of rotation angles in parameterized gates, admittedly at the cost of expressibility. Recent work on quantum feature maps and kernel methods has found links to the classical theory of kernel, implying that trainability in quantum machine learning can be informed by work on the large canonical literature on kernel methods and support vector machines. It is still unclear whether quantum advantage was reached in machine learning and classical algorithms have shown significant performance increases as well as theoretical heterogeneous indications of constraints on quantum speedups on particular learning tasks.

Many-body quantum systems Quantum simulation of many-body quantum systems is among the earliest applications in which quantum computers can be of exponential benefit compared to classical methods in undertaking certain tasks. Variational quantum simulation Preparation of approximate ground states or thermal states of many-body Hamiltonians, dynamics of many-body Hamiltonians or non-trivial dynamics, or correlation functions and response properties, can be computed with parameterized circuits. Physical system and target system determine the trainability issues of quantum simulation. Local interaction systems and area-law entanglement systems which can be in one-dimensional quantum chains or in weakly interacting systems tend to be more easily trainable than systems with long-range interactions or volume-law entanglement. The process of quantum simulation time evolution comes with special trainability concerns because the length of a circuit needed to simulate time accurately (generally) increases with the length of the simulation time, and may cause barren plateaus of the long-time dynamics.

The specialised simulation tasks such as quantum phase transition studies, thermalization dynamics and transport phenomena also possess certain trainability considerations. Examples where trainability is most hard to achieve include near quantum phase transitions when the correlation lengths are very different, and entanglement is increasing very quickly, because of the critical fluctuations and long spectral correlations. Dissipative dynamics and open quantum systems simulation adds

further complexity due to coupling with environments and effects of decoherence which have to be included in the variational ansatz. Recent developments have seen the trainability of types of quantum simulation through the use of tensor-network-based ansatzes that inherently capture the structure of entanglement (symmetry-adapted circuits), and sequential simulation protocols that prepare complex states sequentially out of simple initial states (sequential simulation). The analysis of deck forces and moments based on simulation applications error mitigation has allowed an accurate physical observable to be extracted even out of noisy quantum hardware.

4.6 Hardware Considerations and Noise Effects

The fact of the noisy intermediate scale quantum devices also significantly costs the realistic landscape of variational quantum algorithm implementation, where interactions between hardware flaws, trainability, and optimization are complicated causing way beyond the theoretical abstractions. State of the art quantum processors have a myriad of sources of errors such as faulty implementation of gates, measurement errors, decoherence during processing, qubit-qubit cross-talk, and time dependent calibration errors. Such hardware limitations do not only worsen the quality of the results computed, but also change the nature of the optimization landscape being traversed in quite unexpected manners both exacerbating and sometimes mitigating the difficulties of trainability. The behaviour of such hardware effects and the associated strategies to sustain VQA performance behaviour in them forms a very important frontier towards the practical computation advantage of quantum computing.

One of the pillars of any hardware constraints on VQA trainability and performance is the gate fidelity limitations. Single-qubit gate fidelities of current quantum processors are between 99.9% and 99.99% and two-qubit gate fidelities are between 95% and 99.5% based on the hardware platform, qubit modality and the implemented particular gate. Though these values can be impressive, these errors of the gates at a deep quantum circuit can be lethal in total effect. With one hundred and two qubit gates in a circuit with each gate having a fidelity of 99% the frequency of execution of a circuit with one gate is three-seventh. That is that with one hundred and two qubit gates in a circuit there will be only 37 chances to succeed a circuit with one gate free of error. This error accumulation directly imposes a direct impact on VQA optimization, adding noise to an evaluation of cost functions and gradient estimates, is more or less an additional stochastic element into the optimization landscape, which is structure- and depth-dependent.

The noise-corrected cost function optimised by VQAs in actual hardware does not match the original noiseless cost function in a significant way which can change the optimization problem fundamentally. Depolarizing noise, where the single intended

quantum state is randomly depolarized to the maximally mixed state with some probability, has the propensity of drawing expectation values to the maximally mixed state values of the observable, and tends to be zero in the case of traceless observables. The impact of this effect forms noisy local minima in which the noise-compromised cost function is optimised, and a noiseless problem has not been solved. Worse, noise can distort the gradient landscape to alter the magnitude and direction of the gradients in a manner that is inaccurate by affecting optimization towards the real solutions. Theoretical literature in recent years has described such noise induced spatial landscape changes in the context of the different noise models and found it to explain that some sets of coherent error can in fact fracture barren plateaus through the reduction in the depth of effective circuits at the expense of biasing solutions towards the correct solutions.

Another constraint that is felt by the trainability limitations on VQA circuit design is hardware topology and connectivity component constraints of qubits. In the majority of quantum processors today, connectivity is limited with each qubit being able to interact with a very limited number of neighbours usually ranging between two and six qubits depending on the architecture. The gates between non-adjacent qubits cannot be implemented without any swap operations or additional routing operations which add depth and additional error to circuits. This limitation of connectivity makes tradeoffs between hardware-native sets of gates with limited connectivity and high fidelity and circuit compilation to realise desired gates with higher depth and error. In the case of VQA trainability, such considerations dictate the watchfulness of the circuit design to hardware-efficient circuits that do not disturb native connectivity but have no correspondence to the structure of problems and may train poorly compared to optimally structured ansatze.

The connexion between noises in hardware and barren plateaus provides some complex phenomena as well as counterintuitive developments. On the one hand, noise decoherence may hamper unnecessary increase of entanglements and effectively decrease the magnitude of quantum coherences within the circuit possibly mitigating barren plateaus caused by entanglements. Other studies have found that with some cases of problems, the trainability in moderately noisy circuits is better than noiseless simulations. Nevertheless, this realisation presents a huge price: not only do the decreased entanglement decrease the quantum correlations required to achieve quantum advantage, but the cost as modified by noise can be a misrepresentation of the original problem. Moreover, noise may develop its own type of barren plateaus by mixing towards maximally mixed through noise. This is a delicate scenario where the classification of the noise type, its strength, circuit characteristics and the nature of the problem are to be considered carefully and in each different case.

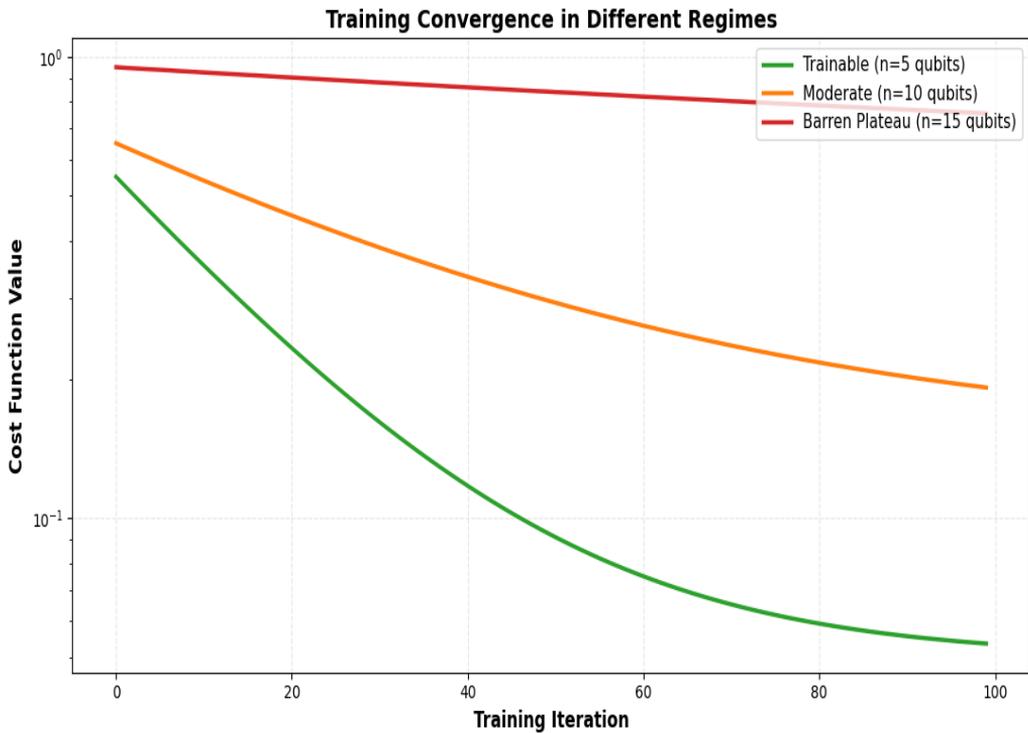


Fig 4: Training Dynamics Comparison

Error mitigation methods are becoming important mechanisms of regulating the quality of VQA on noisy hardware without implementing complete quantum error correction. The purpose of these methods is to obtain precise expectation values on noisy measurements in the form of classical post processing, which is effectively likened to denoising the cost functional values on which optimization takes place. Zero-noise extrapolation purposely alters the level of noise in the quantum circuit by either stretching or folding gates, tests the cost function at various noise levels, and extrapolates to approximate the value of the extrapolated quantum circuit point at zero noise levels. Probabilistic error cancellation Represents ideals gates as a linear combination of noisy gates with the help of a quasi-probability representation of quantum gates, and constructs unbiased estimators by classically post-processing the measurements of different noisy circuits. Symmetry checks and subspace techniques harness knowledge of known properties of the target state e.g. a quantum number or physical constraint to used in projecting results of noisy measurements onto physically legal subspaces.

Error mitigation methods are much more effective and applicable when used with specific VQA applications and hardware platforms. Zero-noise extrapolation uses more quantum resources that it would otherwise need by running circuits on artificially enhanced noise levels which can lead to a lower signal-to-noise ratio. The method is best when the noise is moderate, and when the noise has a smooth scaling, however it may

not be the most effective with highly noisy systems, and non-Markovian noise. Probabilistic cancellation of error only In theory, probabilistic mitigation of error is theoretically capable of implementing perfect cancellation of error, but in practise grows exponentially with the strength of noise and cannot be effectively implemented. Usage Recent results on learning-based error mitigation deploying machine learning to learn noise models based on calibration data and create approximate error mitigation protocols with a polynomial overhead are promising for the ability to scale to larger systems. Bringing together error reduction and VQA optimization, as a new group of methods, such as gradient error mitigation and error mitigation-sensitive training, form a research direction that includes important opportunities to enhance the real-world performance.

Measurement error mitigation is a method that is directly in response to imperfect readout of quantum states that influences the evaluation of all VQA cost functions. Measurement errors, which are usually typified by the likelihood of measuring the quantum state incorrectly as $|0\rangle$ when actually it is $|1\rangle$ and the reverse, can be combated and reversed using measurement error mitigation methods. Through measuring calibration states by known preparation, constructing a confusion matrix to describe measurement errors, and combining measured powers of the aid states of unknown quantum state to build a sort of realistic unbiased estimators of the true probabilities of the quantum state. The reason why this is a particularly successful method is that measurement errors are frequently the units of error leading to performance in the present-day quantum processors, combined with single-qubit measurement fidelities that are commonly lower than the fidelities of the gates being implemented. Future developments, such as tensored measurement mitigation of multi-qubit observables and continuous recalibration procedures to deal with the effect of time have also helped enhance practice performance even more.

Quantum processors have an additional temporal stability factor and calibration drift, which implies new practical constraints to VQA implementation: for any algorithm that needs thousands of circuit executions in hours or days to optimize, the temporal stability and calibration drift becomes a real issue. Quantum processors are various features such as qubit frequencies, gate parameters, and noise rates, which change with time with the effect of environmental variations, flux noise in superconducting systems, charge noise in semiconductor devices, and other impairments. Such a drift is essentially modifying the optimality of the optimization process throughout the optimization process, which may nullify the calculations of gradients or cost functions in the past. Techniques of calibration drift management comprise regular recalibration of gates and measurements, grouping quantum circuit executions (also of measurements) over time intervals over which the processor remains stable, and classical VQA monitoring of processor drift to identify and modify parameters over time. Other VQA implementations have made use

of online learning strategies that maintain noise models continuously and adjust optimization policies to keep the performance steady even in the case of hardware drift.

The non uniformity with respect to various qubits or gates in a quantum processor makes VQA circuit design and optimization further complicated. All qubits do not all have faithfully coherence times, gate fidelities and measurement-precision, and differences of factors of two are frequently observed within one device. In a similar manner, two-qubit gates between qubits pairs have a significantly different error with respect to the strength of coupling, cross-talk, and precision of control. This heterogeneity implies that the choice of qubits and circuit mapping may affect the performance of VQA greatly, with problem qubits of interest preferably getting implemented upon the best physical qubits and critical gates using the most guaranteed qubit pairs. In recent work on noisy compiled circuit design and VQA circuit design designs, the role of this heterogeneity is specifically considered to choose how many qubits to use and when to schedule the gates to minimize the expected errors without compromising circuit functionality or trainability.

4.7 Emerging Techniques and Future Directions

The high rate of variational quantum computing development has pioneered many new methods that further the limit of what is possible using the existing and upcoming quantum machines and at the same time suggests future research directions that could ultimately revolutionise the field. These are new directions in theory offering new insights into trainability and optimization and in algorithm design providing new training paradigm and circuit architecture, hardware developments to support new capabilities and performance levels, and cross-disciplinary intersections bringing zero-classical machine learning, optimization theory and statistical physics. The bottom line is that it is necessary to understand these new directions and how they may affect VQA research in order to predict the overall trend of future development and what avenues might be defined as the most promising ones towards the attainment of practical quantum advantage.

A quantum-conscious meta-learning is one of the most promising new models, which may lead to a radical advancement in VQA performance by training useful optimization schemes on experience across families of problems. In contrast to classical machine learning in which the meta-learners learn increasingly inductive biases that are typically useful, quantum meta-learning has to overcome quintessence quantum measurement limitations, exponentially large state spaces, and the phenomena of quantum and classical interaction. Recent methods have revealed that meta-learned optimizers are able to discover parameter initialisation strategies, adaptive learning rate schedules, as well as circuit structure alterations that achieve great benefit over hand-crafted ones. The

meta-learner is an architecture usually modelled as an instance of a classical neural network or transformer, which monitors the sequence of values of cost functions, gradients, and circuit weights throughout the process of optimising a training problem, but then makes predictions of the optimal next steps on new problem instances.

The theoretical principles of quantum meta-learning make a connexion to the classical meta-learning theory and explain the unique quantum phenomena. The generalisation bounds of meta-learned VQA optimizers should take into account the classical generalisation to problem distributions on the one hand, and the quantum complexity of the evaluation of cost functions and gradients on the other hand. Recent activity has proven the bounds of PAC-learning style, quantum meta-learning, showing the number of training examples and training problems needed to learn optimizers which are applicable to new training examples. These conceptual findings indicate intriguing trade offs between the variety of training issues, quantum resources that can be used on each training issue, and the generalisation performance that can be obtained. Concerning the best architecture of quantum meta-learners, the limit of using the quantum structure to take advantage of meta-learning beyond the classical techniques, and whether quantum benefit in meta-learning can offer quantum benefit in the downstream VQA tasks are open questions.

The differentiable synthesis of quantum circuits and architecture Searching Differentiable synthesis and architecture search are new technologies in automating VQA circuit design by the use of gradient-based optimization or reinforcement learning. Instead of directly designing an ansatz structure by intuition in a problem, the methods of these techniques can be viewed as optimising a variable which is the circuit architecture, and search the circuit architecture space to find designs which are particularly optimal on various criteria such as expressibility, trainability, hardware efficiency, and problem-specific performance. Classical deep learning methods of neural architecture search serve as conceptual guides, whereas quantum circuit search has an array of different problems such as discrete search space, high-cost evaluation of candidate architectures, and the fact that the search requires quantum coherence requirements across the circuit. More recent applications have used differentiable relaxations of discrete architecture selections, evolutionary architectures with co-evolution between quantum and classical, and agents with reinforcement learning that penalize circuit synthesis by using discrete evolution and selecting a sequence of gates by sampling among a pool of available gates.

The combination of classical and quantum machine learning in hybrid models is another line of advancement with a great potential of breaking the limitations of each paradigm separately. Harolding hybrid architectures allow greater use of classical models through dimensionality reduction and feature extraction, quantum models with quantum processing in areas that quantum advantage might be demonstrated and which therefore

might use quantum computing, and classical post-processing to give a final prediction or decision. In recent years, there has been quantum layers in classical neural networks, where the gradient is propagated through quantum and classical subsystems with back propagation, quantum kernel methods, which use classical kernel learning, to achieve better expressiveness, and alternating quantum-classical networks, in which information is processed by each side of the network in turn. The hybrid methods are possibly in a position to enjoy the advantages of the complementary advantages of classical and quantum computation and reduce the limitations of each one.

The quantum natural gradient and second-order optimization methods have also become useful alternatives to the popular gradient descent due to geometric structure in the quantum state manifold and converting this structure to better performance. The quantum geometric specific of the generalised Fisher information metric of quantum systems, it gives a natural Riemannian metric on the quantum state space parameterized by VQA circuits. Natural gradient optimization in this geometrical form can be used to effectively do steepest descent on the quantum state manifold instead of the parameter space, which can remove underconditioned landscapes and can be performed more quickly. In practical applications, the quantum geometric tensor should be estimated by quantum measurements, which, conventionally, involves circuit executions but current technological progress in the estimation of quantum measurements has made this unwarranted in light of near-optimal geometric tensors estimations such as randomised measurements, classical shadows as well as approximate geometric tensors.

Error mitigation and noise-adaptive training have become the next important directions of VQA expansion even in the conditions of realistic hardware noise. In addition to conventional methods of standard error mitigation as a post-processing, new methods can be found which place the mitigation process at the centre of the training loop by learning noise-sensitive cost functions, learning mitigated gradients, and compile adaptive circuitry which learns noise characterization in real-time. Circuit structure dynamically modified through noise-adapted ansatzes may conceivably preserve both trainability and solution quality through hardware imperfections if all error rates and coherence times are monitored. Error mitigation based on learning builds machine learning models on calibration data and then predicts and removes the effect of noise, potentially leading to higher performance than analytical mitigation models of complicated correlated noise models. Theoretical analysis of these methods is intended to find out when and how mitigation can be demonstrated to enhance the performance of VQA, with respect to issues of sample complexity, the overhead of mitigation, and the accuracy that can be attained.

Compilation and transpilation methods that are optimised to compile QVCA codes, in particular, is a new overlapping field between quantum algorithms and quantum hardware control. VQA-aware compilation has the ability to take advantage of

the iterative properties of variational optimization of a problem, unlike general-purpose circuit compilation, which tries to implement arbitrary circuits as closely as possible. More easily compilable versions of the computation, which tend to either be smoother (as in approximate compilation) or lower error rates (in the form of approximate verification), can actually lead to better quality of VQA on the whole, as it is trainable and can be regarded as sufficiently reliable. Compilation-conscious design Ansatz design Compilation-conscious ansatz design involves the precise design of targets The native gate set and connectivity of the target hardware are taken into account early in the design of the circuit, designing circuits which have a minimal compilation cost. The recent past has seen learned compilers which apply machine learning to find efficient circuits, just-in-time compilation which optimises circuits based on the current hardware calibration, and co-design which achieves both a better structure and a compiler at the same time.

The progress of the theoretical reference points and complexity level definitions of VQA trainability is important preliminary work on the capabilities and limitations of fundamental capabilities. Although barren plateaus have been researched widely, theoretical descriptions of trainability in various classes of circuits, problem classes and in noise models are still lacking. It would be important to create strict complexity separations between trainable and untrainable VQA examples in order to offer invaluable advice in algorithm design and application selection. The question of quantum advantage variational algorithms, such as the ability of efficiently trainable VQAs to solve classical hard problems and in what cases such advantage can be beneficially traded off as in noisy implementation, is open. More recent theory has now started to answer these questions with quantum complexity theory tools, demonstrating lower bounds on the complexity of VQA training on particular types of problems, and has linked this to cryptographic hardness assumptions.

Fault-tolerant variability quantum computing is a further vision in which VQAs are run on error-corrected quantum computers at significantly different resource tradeoffs and considerations. Although existing VQAs are optimised with NISQ machines with superficial circuits and low coherence, with fault-tolerance of the future, they might use much deeper circuits with more complicated entanglement structures, with a potential to provide vastly better solutions and support new uses. The resource demands to run quantum error correction, however, such as large qubit overhead and many error correction cycles beg questions about whether variational algorithm implementation is still beneficial in this regime relative to fully fault tolerant algorithms application. Theory and simulations of fault tolerant VQAs have started to map optimal circuit depths, entanglement structures and trainability of error corrected regimes and new interesting tradeoffs between circuit complexity, the error correction overhead, and the accuracy of the solution have been identified.

Another way in which hybrid algorithmic approaches are going involves combining VQAs with other quantum algorithmic models such as quantum walks, quantum annealing, and quantum-inspired classical algorithms. VQAs based on quantum walks are parameterized quantum walks on graphs as optimization and machine learning networks, whose strengths are potentially offered to graph problems and better trainability characteristics through controllable mixing behaviours. Variational quantum annealing is quantum annealing with the flexibility of variational optimization, where the classical parameters that determine the schedules or intermediate Hamiltonians of an annealing are controlled by classical parameters. VQAs may be used together with quantum-inspired classical algorithms that model some of the behaviour of a quantum computer on a classical computer as a classical pre-training algorithm, a class of verification algorithms, a hybrid algorithm that allocates resources between a classical and a quantum computer based on problem characteristics and capabilities. Such cross-paradigm blends may perhaps take advantage of the strengths of each approach as the weakness of each is alleviated.

Table 2: Challenges, Opportunities, and Future Directions in VQA Trainability

Sr. No.	Challenge Category	Specific Challenge	Current Limitations	Proposed Solutions	Opportunities	Impact Potential	Future Research Directions
1	Barren Plateaus	Exponential Gradient Vanishing	Untrainable Deep Circuits, Global Cost Functions	Local Observables, Correlation-Aware Ansätze	Problem-Specific Circuit Design	High - Enables Scaling	Rigorous Characterization of Trainable Circuit Families
2	Circuit Depth	Decoherence in Deep Circuits	Limited Expressibility vs Trainability Tradeoff	Adaptive Depth, Layerwise Training	Dynamic Circuit Construction	High - Balances Quality and Feasibility	Optimal Depth Determination Algorithms
3	Noise Sensitivity	Hardware Errors Accumulate	Cost Function Bias, Gradient Errors	Error Mitigation, Noise-Adaptive Training	Robust Algorithm Design	Very High - Practical Deployment	Noise-Aware Optimization Theory
4	Measurement Overhead	Exponential Sampling Complexity	Resource Constraints, Long Runtimes	Importance Sampling, Classical Shadows	Efficient Resource Allocation	Medium - Improves Efficiency	Sample-Optimal Gradient Estimation

5	Ansatz Design	Problem-Specific vs Universal Tradeoff	Manual Design Requires Expertise	Automated Architecture Search	Generalizable Ansätze	High - Broadens Applicability	Meta-Learning for Ansatz Discovery
6	Local Minima	Suboptimal Solutions	Poor Initialization Sensitivity	Multi-Start, Basin Hopping	Landscape Engineering	Medium - Improves Reliability	Theoretical Guarantees on Landscape Properties
7	Expressibility	Limited State Space Coverage	Insufficient Circuit Complexity	Deeper Circuits, Novel Gate Sets	Enhanced Computational Power	High - Expands Problem Classes	Expressibility-Trainability Unified Frameworks
8	Hardware Connectivity	Limited Qubit Interactions	Compilation Overhead, SWAP Gates	Topology-Aware Design, Hardware-Efficient Ansätze	Native Circuit Implementations	Medium - Hardware Utilization	Co-Design of Algorithms and Hardware
9	Scalability	Exponential Resource Growth	Problem Size Limitations	Divide-and-Conquer, Hierarchical Methods	Large-Scale Applications	Very High - Practical Impact	Provably Scalable VQA Frameworks
10	Optimization Algorithm	Slow Convergence, Poor Solutions	Classical Optimizer Limitations	Quantum Natural Gradient, Second-Order Methods	Faster Training	High - Reduces Time-to-Solution	Quantum-Classical Co-Optimized Methods
11	Initialization Strategy	Random Start in Barren Plateaus	Wasted Optimization Effort	Classical Pre-Training, Identity Initialization	Improved Starting Points	Medium - Enhances Convergence	Provably Good Initialization Schemes
12	Entanglement Structure	Volume-Law Induced Trainability Loss	Uncontrolled Entanglement Growth	Controlled Entanglement Ansätze	Structured State Preparation	High - Maintains Trainability	Entanglement-Aware Circuit Design
13	Cost Function Design	Global Observables Cause Plateaus	Observable Choice	Hierarchical Cost Functions, Local	Trainable Problem Encodings	High - Fundamental	Observable Decompo

			Constrai nts	Decompo sitions		Design Principle	sition Theory
1 4	Quantum- Classical Interface	Communi cation Bottlene ck	Hybrid Loop Inefficie ncy	Parallel Execution , Batch Processin g	Reduced Latency	Medium - Operational Efficiency	Optimized Hybrid Architectu res
1 5	Benchmark ing	No Standard Metrics	Incompa rable Results Across Studies	Unified Benchmark Suites, Standardi zed Metrics	Rigorous Performan ce Comparis on	Medium - Communi ty Standards	Comprehe nsive Benchmark Developm ent
1 6	Theoretic al Understan ding	Incomplet e Trainabili ty Theory	Unpredic table Performa nce	Geometri c Analysis, Complexi ty Theory	Principled Algorithm Design	Very High - Foundational Knowledge	Unified Trainabilit y Theory
1 7	Applicati on Specificity	Domain Knowledge Requirements	Limited Accessibi lity	Automate d Problem Encoding	Broader Adoption	Medium - User Accessibility	Problem- Agnostic VQA Framework s
1 8	Verificati on	Solution Quality Uncertain	No Classical Ground Truth	Partial Verificati on, Statistical Validatio n	Confidenc e in Results	High - Trustwort hiness	Verificati on Protocol Developm ent
1 9	Error Mitigatio n Overhead	Increased Resource Requirements	Limited Benefit vs Cost	Learning- Based Mitigatio n, Adaptive Schemes	Practical Noise Tolerance	High - NISQ Viability	Low- Overhead Mitigation Technique s
2 0	Hyperpar ameter Tuning	Many Tunable Parameters	Trial- and- Error Optimiza tion	Automate d Hyperpar ameter Optimizat ion	Robust Performan ce	Medium - Ease of Use	Meta- Learned Hyperpara meter Strategies
2 1	Quantum Advantage	Uncertain Advantage Regime	Classical Competition	Complexi ty- Theoretic Analysis, Empirical Demonstr ations	Computati onal Supremac y	Very High - Field Validatio n	Rigorous Advantage Characteri zation

2 2	Heterogeneous Hardware	Device-Specific Optimization	Portability Issues	Hardware-Agnostic Abstractions	Cross-Platform Deployment	Medium - Flexibility	Universal VQA Frameworks
2 3	Real-Time Constraints	Industrial Application Requirements	Slow Optimization Cycles	Amortized Training, Transfer Learning	Practical Deployments	High - Industry Adoption	Real-Time VQA Protocols
2 4	Multi-Objective Optimization	Conflicting Objectives	Pareto Front Exploration	Multi-Objective Evolutionary Algorithms	Complex Problem Solving	Medium - Rich Solutions	Quantum Multi-Objective Optimization
2 5	Symmetry Exploitation	Manual Symmetry Identification	Missed Optimization Opportunities	Automated Symmetry Detection	Reduced Problem Complexity	Medium - Efficiency Gains	Symmetry-Aware Automated Design
2 6	Gradient Estimation	Parameter-Shift Rule Overhead	2 Evaluations Per Parameter	Simultaneous Perturbation, Finite Differences	Resource Reduction	Medium - Measurement Savings	Optimal Gradient Estimation Protocols
2 7	Circuit Compilation	Suboptimal Transpilation	Fidelity Loss, Depth Increase	VQA-Aware Compilation	Native Implementation Efficiency	Medium - Error Reduction	Co-Optimization of Circuits and Compilation
2 8	Overfitting	Limited Training Data	Poor Generalization	Regularization, Data Augmentation	Robust Learning	Medium - QML Applications	Quantum Regularization Theory
2 9	Catastrophic Interference	Sequential Learning Degradation	Forgetting Previous Tasks	Continual Learning, Elastic Weight Consolidation	Lifelong Learning Systems	Low - Specialized Applications	Quantum Continual Learning
3 0	Interpretability	Black-Box Optimization	Unclear Failure Modes	Visualization Tools, Explainable Quantum AI	Understanding and Trust	Medium - Scientific Insight	

5. Conclusion

The whole exploration of trainability and optimization space of Variational Quantum Algorithms discussed in this chapter unveils one of the areas that are on the edge of a knife, between the remarkable theoretical progress and insurmountable practical issues. VQAs have become the most practical route to the derivation of computational value out of existing and near future quantum devices and have proven to be successful in a wide range of different application fields including quantum chemistry and materials science as well as machine learning and combinatorial optimization. The hybrid quantum-classical nature of VQAs enables the circumvention of some of the most challenging tasks in pure quantum computation, such as the need to build deep circuits and complete quantum error correction, and an abundance of the benefits of quantum and classical computation. Nevertheless, the trainability problems that are exemplified by the barren plateau phenomenon are intrinsic ones that are based on the quantum mechanical features of high-dimensional Hilbert spaces, but not on the real implementation issues.

Theories that have been developed in recent years have served as fundamental understanding of the processes involved in trainability loss with complex relationships between entanglement structure, cost functional locality, circuit architecture, and gradient scaling. The description of barren plateaus by the use of concentration of measure, quantum Fisher information geometry and Lie algebraic analysis has worked wonders in changing our perception of the world based on empirical data, to one based on formal mathematical models. These theoretical discoveries have facilitated the creation of the principled mitigation techniques such as correlation-aware ansatz, layerwise training protocols, quantum natural gradient methods and problem-specific circuit designs that utilise the symmetry and physical structure. However, there are still considerable gaps in theory, especially within expressibility-trainability trade-off issues, the effect of realistic noisy models on optimization behavior, as well as the nature of ultimate limitations to the dynamics performed by variational methods.

Practical implementation of VQAs on modern quantum devices adds further complexity due to the effects of noise, low connectivity, calibration drift and restricted measurement budgets. Although error mitigation methods have proven to be incredibly successful in extracting correct results using the noisy quantum processors, the overheads of the methods and their limitations to larger systems are also unknown. The complex interaction of hardware flaws with trainability, in which noise can furthermore both decrease barren plateaus and at the same time bias solutions and generate spurious local minima, yields a hard-to-navigate optimization problem, necessitating advanced navigation methods. The variability of quantum hardware as various platforms have their own benefits and constraints also makes it harder to create general VQA solutions.

In the future, some of these potential research directions hold some promising prospects as to possibly eliminating current shortcomings and scaling variational quantum algorithms to be trainable. The combination of machine learning methods such as meta-learning, automated architecture search, and learning-based error mitigation is especially promising to find the most competitive optimization solutions and circuit design proposed by hand. Improving the rates of convergence and solution quality could be dramatically improved by designing quantum-mindful classical optimizers based on natural gradient algorithms with geometric structures of quantum state spaces, as well as second-order algorithms. The problem structure is understood and used in computational architecture design with advancements in knowledge of symmetries, locality and physical constraints, which can be explored to design the structure of an ansatz that remains trainable and provides the desirable expressibility. The discussion of other training paradigms such as the use of reinforcement learning in circuit assembly, evolutionary strategies that optimise problems in a global approach or the use of hybrid quantum-classical models which assigns computational tasks according to comparative advantage can open up new opportunities.

The road to the practise of quantum advantage with VQAs needs further developments on various fronts. Theoretical progress should define strict characterizations of trainable circuit families, create a unified frameworks relating expressibility and trainability as well as establish fundamental limitations on attainable performance on realistic constraints. The innovations in algorithms have to generate scalable training procedures, commonly resilient methods of diminishing errors, and domain-specific optimizations. Other hardware enhancements such as improved gate fidelities, longer coherence time, higher numbers of qubits and improved connectivity will enlarge the problem space accessible. Interdisciplinary teamwork among quantum physicists, computer researchers, optimization scholars, and domain knowledge workers will play a crucial role in the translation of the algorithmic skills into practise and generating real value through them.

It remains yet to be seen of VQAs whether practical quantum computational advantage will be realised soon, depending upon a fresh theoretical contribution, algorithmic creativity, and hardware breakthrough. The trainability issues reported in this chapter are daunting and not insurmountable and new mitigation methods are showing that with a careful design of an algorithm, it is possible to continue to maintain a optimization possibility even with moderately large systems. With the ongoing advance in quantum hardware and further theoretical development, the space of problems with which VQAs can outperform classical models is probably going to grow, and this may include industrially important problems in drug discovery, materials design, financial optimization, and machine learning. The overall theory on the understanding of transitioning VQA trainability represented in this chapter, not only offers a baseline on

the ongoing research but also provides a prospective guide on how the field will avail of new opportunities and overcome the remaining challenges on the way to practical quantum advantage. VQAs have become the most promising state of this PCA to computational value extraction of existing and near-term quantum devices and are proving to be successful in a variety of application fields and areas, ranging in scope and utility, between quantum chemistry and materials science to machine learning and combinatorial optimization. VQAs have a hybrid quantum-classical architecture, which gracefully avoids some of the most formidable challenges of pure quantum computation such as the demand of deep circuits and full quantum error correction, and exploiting the more complementary capabilities between quantum and classical computation. Nevertheless, the nature of the trainability struggles embodied by the barren plateau effect are the essential complications connected to the quantum mechanical nature of high dimensional Hilbert spaces and are not simply the implementation issues.

The discussion of the theoretical framework that has been developed recently has been instrumental in the support of the mechanisms underlying the training degradation, that entangled entanglement structure and the locality of cost functions, circuit structure, and gradient scaling. The analysis of barren plateaus in terms of the concentration of measure, quantum Fisher information geometry, and Lie algebraic geometry has changed our perception of the phenomenon based on the empirical results to the mathematical rigorous constructs. Theoretical progress has made possible principled mitigation techniques such as correlation-aware ansatzes, layer-based training procedures, quantum natural gradients as well as problem-specific circuit designs that make use of symmetries and physics. However, there are still considerable gaps in theory, especially within expressibility-trainability trade-off issues, the effect of realistic noisy models on optimization behavior, as well as the nature of ultimate limitations to the dynamics performed by variational methods.

The implementation process of VQAs in state-of-the-art quantum hardware introduces further complexity domains to them by the impact of noise effects, network connectivity, drift in calibration, and finite resources in the determination of measurements. Although methods of mitigation of errors have proven to be extremely effective in obtaining the correct results to noisy quantum processors, the overhead of error reduction schemes to scale to more complex systems and the ability to scale the technique to larger scale remain unclear. The complex interaction between hardware flaws and trainability, with noises being able to both correct barren plateaus and also introduce bias in solutions, and to cause spurious local minima, provides a difficult optimization problem that may need complicated navigation strategies. The heterogeneity of quantum hardware where each platform has its own advantages and disadvantages also makes the creation of universal solutions to VQA even more problematic.

Moving forward in time, there are a few promising directions of research that provide avenues to a solution of current limitations and scalable, trainable variational quantum algorithms. The combination of machine learning methods such as meta-learning, automated architecture search, and learning based error mitigation has certain potential in finding optimization strategies and circuit designs which can be superior to hand-designed one. Improving the rates of convergence and solution quality could be dramatically improved by designing quantum-mindful classical optimizers based on natural gradient algorithms with geometric structures of quantum state spaces, as well as second-order algorithms. Discoveries in the field of recognising and using structure of problems via symmetries, locality and physical constraints are some methods that may be used to design ANsätze that preserve trainability and reach required expressibility. The new possibilities might be unlocked by the exploration of new training paradigms such as reinforcement learning to construct circuits, evolution to optimise the entire problem globally, and hybrid quantum-classical models which assign computational work in accordance with comparative advantage.

The road to the realistic quantum advantage with the help of VQAs still needs advancement in several directions. Theory Diverging theoretical benefits have to define rigorous versions of families of trainable circuits, construct coherent structures between expressibility and trainability, establish basic constraints on performance under realistic conditions. The algorithmic innovations should generate a scalable training protocol, error mitigation strategies, and domain-specific optimizations, which take advantage of the domain knowledge. The enhancements in hardware such as improving gate fidelities, longer coherence times, larger qubit counts, and improved connexion will give access to larger problem space. The cross-disciplinary interaction with quantum physicists, computer scientists, optimization researchers, and domain experts will be necessary in the translation of algorithmic capabilities to practical uses, which have provided true value.

The eventual realisation of useful quantum computational advantage by VQAs is an unresolved problem, which depends on further theoretical development, further algorithm development, and hardware development. These trainability obstacles reported in this chapter are not trivial, but solvable, and the emergence of mitigation measures shows that optimization is still maintainable in even moderately sizable systems by careful design of their algorithms. With the continued improvement of quantum hardware and our theoretical understanding of it, the regime of problems within which VQAs will be able to outperform classical methods may continue to grow and possibly include applications of relevance to industries such as drug discovery, materials design, financial optimization, and machine learning.

Chapter 4: Quantum Neural Network Architectures: Depth, Width, and Expressive Power

1 Abstract

Quantum neural networks (QNNs) constitute a paradigm shift convergence between quantum computing and machine learning, and they present unexplored computational advantages by harnessing the effects of quantum mechanics, i.e., superposition, entanglement or interference, etc. In the chapter, a detailed analysis of quantum neural network structures has been given with a special focus on the structural parameters that drive the depth in the architecture, the width, and the combined expressiveness. The focus of the investigation includes the theoretical bases that separate quantum architectures using classical ones, and discusses how quantum gates, variational circuits, and quantum feature map interactions can be used to make models with the ability to model high-dimensional, complex functions. This chapter summarises the novel trends in the principles of QNN design, such as hybrid quantum-classical architecture and barren plateau mitigation techniques and the principles of novel entanglement-based connectivity patterns, through systematic literature review using the PRISMA methodology. The discussion focuses on such vital issues as circuit trainability, quantum noise resilience, and scalability limits and determines quantum advantage transformative opportunities in particular areas of problems. The correlation between architectural decisions and expressibility measurements is paid special attention and the depth-width trade-offs effect that can affect the representational capacity and practical implementability relevance on near-term quantum space machines are addressed. The chapter includes both thorough results of simulation of quantum chemistry, optimization studies, generative modelling, pattern recognition, as well as an in-depth discussion of newer methods such as quantum attention models, dynamically configured circuits, and error-resilient training algorithms. Incorporating both theoretical progress and experimental results on the implementations of quantum hardware systems as at today, this paper provides a comprehensive framework of how quantum neural networks are

understood and how the current generation of noisy intermediate-scale quantum computing can be understood to be developed.

2. Introduction

The collision of quantum computing with the neural network theory is one of the most radical advancements in the field of computational science, which is posed to transform the way we solve the challenges of recognising complex patterns, optimal solutions, and learning tasks. Quantum neural networks arise due to the discovery that quantum systems, by virtue of their superposition capabilities and their formation of complex entanglement structures may be computationally useful in particular categories of machine learning problems that are not solvable by classical systems. The quantum neural networks differ (possibly radically) in the aspects of their architectural design than the classical network architectures: the building blocks and mechanism of quantum neural networks are based on quantum mechanics instead of on a regular digital computational platform. In classical neural networks, layers of neurons are represented by edges with weights between them, quantum neural networks are made of chains or circuits of quantum gates on qubit registers, and parameterized rotation of qubit registers play similar functions to trainable weights. The notions of depth and width, though loosely analogous to those of classical computation, take on different values in the quantum theory: and depth is how many consecutive steps of the circuit are performed, and width is how many qubits are used in the computation. These building blocks suggest a direct connexion between these architectural parameters and the expressive power of the network: its ability to approach target functions, though the relationship in quantum neural networks with uniquely quantum phenomena such as entanglement structure, gate fidelity and measurement collapse mediate this relationship. Noisy intermediate-scale quantum (NISQ) devices facilitate experimentalism in quantum neural networks offering tens to hundreds of qubits but without fault tolerance to quantum computation. This technological environment has given work in architectural research variational quantum algorithms, in which shallow and moderately depth circuits with trainable parameters are optimised by hybrid quantum-classical algorithms. VQE and quantum approximate optimization algorithm (QAOA) have provided the principles templates related to the design of QNNs that showed how parameterized quantum circuits could be trained to learn to do particular tasks regardless of the imperfections of the hardware. The theoretical work on the understanding of quantum neural network architectures has led to a vigorous exploration of the factors that control the trainability, repressibility, and realistic performance of quantum neuromorphic engineering on near-term devices. The conceptualization of quantum neural network architectures is based upon quantum information theory, computational learning theory and statistical physics. The quantum circuit complexity theory offers the means of studying the influence of architecture-level

decisions on the computational resources to prepare certain quantum states, or apply other transformations they want. Tensor network representations can provide different ways of understanding quantum circuits, exposing network topology information and entanglement structure information and give ways to classically simulate and analyse quantum circuits. Information based measures such as the quantum Fisher information and entanglement entropy have come to be important metrics of the ability of a circuit to encode and manipulate information, and directly relate the parameters to learning behaviour in architectural terms. Recent discoveries in quantum machine learning have both indicated potential capabilities and limitations in quantum neural network architectures. Conceptual findings have proven quantum advantage of particular learning tasks, showing exponential discrepancies in sample intricacy or calculation effectiveness of issues such as learning particular quantum states, modelling quantum systems, and resolving organised optimization issues. Parallel to this, like barren plateaus, the one-synthetic vanishing gradients of randomly set-up deep quantum circuits, in both architectural design and initialization methods their crucial role has been noted. The finding of these training pathologies has inspired studies into structured architectures, such as quantum convolutional networks, quantum recurrent networks, and tree-tensor networks, may have connectivity constraints that are easier to train while having the same expressive power.

3. Methodology

The methodology used in this chapter is the systematic literature review that relies on the principle of Preferred Reporting Items of Systematic Reviews and Meta-Analyses (PRISMA) with its specific alterations in the context of an interdisciplinary quantum neural networks domain. This review includes peer-reviewed journal articles, conference proceedings, preprints of known archives, and technical reports between 2015 and 2025: the timeframe of intensive research in quantum machine learning, as well as quantum computing of the NISQ era. The systematically search of electronic databases such as IEEE Xplore, ACM Digital Library, arxiv, Springer, Nature Portfolio, the American Physical Society journals, was carried out through the use of controlled vocabulary and key-word combinations such as, quantum neural networks, variational quantum circuits, quantum machine learning architectures, parameterized quantum circuits, quantum repressibility, etc. The first search had brought about 1,847 potentially viable publications that underwent screening depending on relevance in titles and abstracts to the subject of architectural considerations in quantum neural networks. Evaluation of 423 articles was performed in full-text by the use of an inclusion criteria that was limited to original research related to depth, width, power of expression, architecture, or realisation of quantum neural networks. Exclusion criteria were used to remove purely theoretic papers on quantum computing with no component of machine learning,

classical neural network studies with no quantum component and review papers with no novel synthesis or analysis. The information about architectural parameters, theoretical outcomes, or experimental examples, the spheres of its application, quantitative metrics of its performance, and the issues identified was extracted in the data. Quality was measured on the basis of methodological rigour, reproducibility, theoretical soundness and where necessary, empirical validation. Synthesis processes created a synthesis of theoretical, computational and experimental work in building up very complete knowledge of architectural principles, underpinnings, and how they would be applied to quantum neural network design.

4. Results and Discussion

4.1 Applications of Quantum Neural Network Architectures

The quantum neural network architecture application space has increased significantly, as the abilities of quantum hardware and algorithmic development have both increased, and the area of operations where architectural design decisions play a decisive role in the performance and the possibility of quantum advantage [9,24-26]. Probably the most developed and potentially most promising application of quantum neural networks is quantum chemistry, as well as molecular simulation, which is essentially a quantum mechanical problem, and which is exponential in complexity when solved using traditional relying methods. QNN-based variational quantum eigen solvers have already been shown to compute ground state energies of molecular Hamiltonians, and architectural engineering choices such as depth and width do have a direct effect on the accuracy with which an energy can be calculated and the convergence must of the optimization procedure. The hardware-efficient ansatz, which organises quantum circuits to be compatible with the naturally available gate sets and connectivity of particular quantum hardware, is one example of how application demands and architecture design; focusing on whether it is practical to implement and robust to noise instead of what abstract theoretic repressibility can do.

In quantum chemistry novel applications have considered progressively complex molecular systems, extending the limits of what can be modelled with the current quantum hardware and providing an insight into architectural considerations, to achieve a practical quantum advantage. This has spawned the creation of quantum neural network designs with symmetry conditions and conservation laws built directly into the circuit, inspired by the strong level of correlations of electron systems hard to compute with the aid of classical approximation schemes and configuration interaction dynamics. These symmetrized architectures minimise the parameter space with a fixed level of

relevance physics, allowing to optimise much more efficiently and provide a higher level of efficiency with finite quantum resources. The unitary coupled-cluster ansatz and its variants are some of the most notable applications of the presence of chemical knowledge in architectural design, in which parameterized quantum circuits are designed to imbue the creation of the correct classes of entangled state of both molecular ground and excited states. The scaling of depth requirements in these applications is based on molecular size and accuracy of interest, and width is a direct function of the orbitals or number of spin modes modelled in the quantum simulation, which form the basic resource requirements that pose a challenge to current NISQ devices.

Financial calculus and portfolio optimization have become the attractive spheres of application of quantum neural networks especially in areas that have complex relationships, pricing many dimensional options, and the evaluation of risk in uncertainty. Quantum neural network architecture Quantum neural network architectures that are applied to financial tasks typically focus on quantum feature maps to encode classical financial data into quantum states to be acted on through variational circuits to learn correlations and dependencies useful to a prediction or optimization task. Such networks are usually determined by the dimensionality of the financial data that is being processed and depth affects the complex relationships that the model will be able to establish between various assets, market factors, or time dependencies. There has been an investigation of quantum approximate optimization algorithms modified into designer other quantum neural network designs, on portfolio optimization problems, where all aspects of the architecture including mixer and problem Hamiltonians, circuit depth, and choice of measurements have a direct effect on the quality and speed of the solution. The possibility of quantum advantage in such applications is a research question, and the architectural efficiency, the capability to achieve many complex financial relationships with a smaller quantum resource than the straightforward classical methods or classical resources, is an important measure of usefulness in practise.

Another important field of application to quantum neural network architecture is generative modelling, in which the property of quantum amplitude-based representation and manipulation of probability distributions provides unique functionality. Quantum generative adversarial networks (QGANs) use parameterized quantum circuit generators such that the measurement statistics of quantum states generated by the quantum generators matches a set of probability distributions of interest, and in such tasks, discriminator networks (quantum or classical) give training signals. The patentability of the generator architecture, which is dictated by its level of depth and width as well as precision of gates, essentially limits the category of representation that could be educated, with a choice in training pronouncement and convergence characteristics as well. Born machines are another paradigm of generative models that incorporates the Born rule, in which circuit architecture is now based on optimizing to get as many

capabilities as possible within a specified quantum resource budget. Quantum generative models have been used in applications through synthetic data generation, unsupervised training of quantum states, and useful in modelling complex probability distributions occurring in machine learning and statistical physics, with diverse architectural requirements across such applications.

The classification of images and pattern recognition jobs have been studied through the application of quantum neural network designs, but the route to quantum benefit on classical information analysis has not been as evident as on naturally quantum issues. Quantum convolutional neural networks generalise the concept of the classical CNNs to quantum systems, using the local parameterized quantum circuits on the overlapping areas of quantum-encoded images, and then pooling on the circuits can be done by either partial measurements or quantum gate. The architectural analogies to classical convolutional networks offer natural design rules as well as utilise quantum effects such as interference and entanglement to potentially make them much more effective at extracting features. The width of such architectures is equal to the resolution of images and the number of qubits allocated to the representation of pixel values, whereas the depth refers to the number of convolutional layers, as well as the complexity of local circuits. Experimental rendezvous on quantum devices have provided evidence of image classification prototype generated quantities, but practical issues of encoding overhead, circuit depth factor, and the advantageousness of new entrusting the systems of classification onto the *vie in quanta* by the issue of effectively inputting negative image details to quantized states have persisted.

Quantum neural network architecture has been inspired to solve optimization problems associated with a wide variety of areas such as logistics, scheduling, allocating resources, and solving constraint-satisfaction problems. These are frequently implemented in graph based quantum neural networks such that qubits are represented as nodes or edges and that important problem conditions and objective functions are encoded into the structure of parameterized circuits. The level of optimization-oriented QNN architectures is often associated with how many iterations it requires to reach its constraints, and objectives, and width with problem size -number of variables or constraints in the instance of optimization. The group of quantum alternating operator ANNN design architectures are an example of application-specific structure being used to design quantum neural networks, where alternating layers of level-sensitive operators and mixing operators are used, which at all times guarantee feasibility, and particle through the solution space. Architectural decisions made that Favor a trade-off between repressibility and implement ability on near-term hardware combined with the efficiency with which problem structure is encoded into quantum circuits are the key to whether a quantum advantage can be achieved in optimization applications or not.

Other methods that have so far been implemented using quantum recurrent neural network architectures are time series prediction and sequential data processing, which adds temporal dynamics and memory into the design of quantum circuits. These architectures have even more difficulty in preserving coherent quantum states between sequential processing steps and have feedback mechanisms similar to classical recurrent connexions. The Quantum long short-term memory (QLSTM) models simulate memory cell functions on classical LSTMs by using quantum gates and partial measurements to select and discard quantum state at each timestep in a Being constructed of parameterized quantum circuits, the architectures implement the quantum long short-term memory (QLSTM) paradigm that represents quantum long-term memory. Quantum recurrent architectures are characterised by depth in both the recurrent depth (reduction in the number of timesteps) and circuit depth (reduction in the number of gates per timestep), and generate compounding coherence demands putting a strain on existing quantum hardware. It has been applied in financial time series forecasting, natural language processing and in modelling dynamical systems, and architectural research has been done on balancing storage capacity, repressibility, and feasibility of operation as limited by small coherence times.

Quantum reinforcement learning is a new field of application, with quantum neural network designs being policy or value functional approximators on quantum or classical worlds. Trained variational quantum circuits with policy gradient or Q-learning algorithms provide possible benefits regarding adventure in higher dimensional state-action space with the realisation of quantum superposition and in conveying elaborate value functions with the ability of entanglement expression. Quantum engineering Architectural implications in quantum reinforcement learning are how to encode environmental states in quantum numbers, how to compute policy networks between states and the distribution of actions, and how to compute value networks between states and expected returns. It is the width of these architectures, as a factor of state space dimensionality, and the depth is as a factor of the complexity of policies or value functions that can be modelled.

4.1 Techniques in Quantum Neural Network Architecture Design

Efficient quantum neural network architectures are built using an elaborate repertoire of methods that deal with the special issues and prospects that quantum computers create. The circuit of variational quantum quantum circuits forms the basis of the vast majority of the current systems employed in quantum neural networks, as they utilise parameterized quantum gates whose rotation angles act as the trainable parameters similar to weights in classical neural network systems. These variational circuits usually place layers of parameterized single-qubit rotation with fixed entangling gates in

between and apply an organised parameterized ansatz, which can be trained using hybrid quantum-classical training algorithms. Problem-specific ansatz design is an important element of quantum neural network design necessitating the attention of careful planning of the quantum states and transformations to be modelled, the discrete symmetries and structure of the problem space, as well as the connectivity and gate potential of the intended quantum device(s). As an example, hardware-efficient ansatzes represent circuit designs that synthesise circuits to be native gate sets and qubit connectivity of particular quantum processors, spending little to no compiler time, and minimising implementation errors with potentially a trade-off in repressibility (compared to hardware) to implement ability.

In most quantum neural network applications, quantum feature mapping techniques offer the interface between classical data and quantum processing with encoding of classical data into quantum states being a crucial step. Amplitude encoding is one model in which classical data vectors are resourcefully distributed into quantum state amplitudes and offer an exponentially compressible representation, however they need complex state preparation circuits with considerable depth levels that could become overwhelming of the advantages of compressed resources. Basis encoding algorithms views classical bit strings as directly as computational basis states, and has the benefits of being easy to implement at the price of linear qubit scale with data dimensionality. Angle encoding entails known classical characteristics in rotating parameterized quantum gates, which offers a tradeoff between parameterized quantum gates implementation complexity and angle encoding performance. More advanced methods of feature mapping such as quantum kernel embeddings and trainable quantum embeddings use variational circuits to learn suitable transformations on the classical feature space to quantum state space, and the architecture of the encoding circuit is now a design choice, which influences the performance of downstream tasks.

The methods of entanglement structuring have a vital impact on the expressivity of quantum neural network architectures as well as their trainability because entanglement patterns define the correlations which those learners can model as well as the quantum advantage that they may be able to achieve. Brick-layer entangling constructions are regular high-level assignments of two-qubit gates i.e. between pairs of even qubits and pairs of odd qubits, allowing complete depth extension. Architectures inspired by tree-tensor networks use hierarchical entanglement schemes, which resemble the form of a tensor network decomposition, and have a theoretical benefit in being able to represent particular quantum states, as well as having trainability properties. Complete entangling networks which put gates between every pair of qubits can achieve full expressibility but are subject to the needs of circuits with multiplication depth, or with mise-en-scene connectivity, placing two- qubit gates between geometrically adjacent qubits, engineering some expressiveness at the cost of sacrificing depth and winning higher

implementability on near-term machines. The adaptive entanglement technique of evolving the structure of entangling operations over training according to performance measurements or information-theoretic quantities is an emerging technique aiming to optimise the structure of the entanglement as an element of design choice instead of as a fixed design choice.

Gradient estimation algorithms to train quantum neural networks have since developed past naive parameter-shift guidelines to include advanced approaches that can deal with the issues of barren plateaus and measurement statistics. The parameter-shift rule allows computing the exact gradient estimating the gradient without bias but multiple circuit evaluations per parameter are necessary. The stochastic gradient estimation methods use either simultaneous perturbation strategies or finite-differentiated evaluation in order to minimise the number of measurements, sacrificing some degree of accuracy in the name of better scaling with parameter size. The natural gradient algorithms applied to the quantum system are based on the quantum Fisher information matrix to use the geometric properties of the quantum state space to generate an optimization path that is more efficient, with the price of making their Fisher information estimation more expensive to compute. Computation-free optimization methods such as evolutionary algorithms, Bayesian optimization, quantum natural evolution strategies, or gradient-free methods are an alternative to gradient estimate where it is impractical, and architecturally may suggest different depth-width trade-offs as to the different gradient-based methods.

Circuit compilation and optimization algorithms special-purpose quantum hardware to a practical implementation of abstract quantum neural network designs, and are pivotal in the analysis of practical performance and resource aspects of architectural designs. In the provision of methods of gate decomposition, high-level quantum operations are decomposed into schedules of made-in-gate about target hardware, and decomposition intensity has a direct correlation with circuit-wide circuit depth and cumulative errors. Algorithms using circuit synthesis also use template matching, rewriting rules, and numerical optimization to reduce the number of gates in circuits and depth without correspondingly increasing their function. Topology-conscious compilation tools consider connectivity constraints of quantum processors, and use SWAP gates or routing operations in place of direct connexions when a direct connexion is not available, and so may add a considerable amount of circuit depth to architectures that were not architecturally designed with hardware limitations considered. Less precise models Compilation Methods The technique of an approximate compilation trades-offs exact functional correctness of a quantum computation with reduced circuit depth, and has emphasized on methods of applying quantum operations with specified error redundancy in fewer gates - a trade-off especially pertinent when the circuit is an architecture of a deep quantum neural network implementing an unspecified-precision computing model.

Such techniques as noise mitigation and error suppression have now become standard parts of the real world quantum neural network design, and architecture decisions affect the type of errors that arise and how well the mitigation techniques work. The dynamical decoupling methods build sequences of gates that will carefully include error mechanisms which might be averaged, but keep the computational operations, and architecture needs to be changed to size decoupling sequences without adding a great deal of depth. Error extrapolation algorithms and schemes run quantum circuits at many levels of noise and make extrapolation to zero noise limits; to do this, architectural flexibility is needed to artificially increase or decrease noises by adding gates or adjusting the circuit to achieve the noise goal. Probabilistic error cancellation Probabilistic error cancellation methods are quantum operations expressed as linear combinations of noisy implementable operations, and bring the benefit of reducing error by quasi-probability sampling, at the expense of more measurement overheads. Symmetry verification and post-selection algorithms capitalise on known symmetries of quantum neural network architectures that find and reject measurement error outcomes making them effective in the case of problems that have sensible symmetries that are verifiable efficiently.

4.2 Methods for Analyzing and Optimizing Quantum Neural Network Architectures

The quantum mechanical characteristics of optimization and scrutiny of quantum neural networks structures demand particular treatment that considers machine learning aims. The first stage is repressibility analysis, which measures the size of the space of quantum states or transformations which a certain architecture is capable of generating, which gives a basic characterization of representational capacity [27-29]. These approaches normally use statistical quantities of the randomness of the distribution of quantized states produced by random parameterizations of the circuit architecture, contrasting the distribution by the Haar measure of the state space to understand how uniformly the circuit architecture can span the space of possible states of quantized systems. State fidelity probability distribution based repressibility metrics, frame potentials implicit to quantum information theory and volumetric measures of accessible state space can give additional insights into architectural capacity. High repressibility implies that an architecture is capable of modelling a variety of quantum states, usually directly associated with scalability to approximations of target functions or states, but too much expressivity is also a sign of barren plateaus and training hardness.

Another important way of describing the architecture of quantum neural networks is entanglement capability analysis since entanglement is a resource that is uniquely quantum capable of producing correlations previously unachievable in classical systems.

Meyer-Wallach quantified measures of entanglement measure the mean entanglement of any possible bipartition of a system that is quantized by giving a scalar measure that grows with circuit depth and scales depending on the structure of entangling gates. Quantum entropy and mutual information based measures provide a more detailed measurement of the distribution of entanglement among and within different subsystems and their change across the layers of a quantum neural network architecture. Instead of analysing scalar summaries of entanglements analysing their complete distribution of Schmidt coefficients can provide information in finer detail about the entanglement structure that can be correlated with particular representational ability or learning efficiency. The dependence of entanglement ability and learning performance is also a dynamic research topic with some evidence indicating that there must be the correct concentration of entanglement to be used to achieve quantum advantage but too much random entanglement can be counter-trainable.

A key and important question that is answered through trainability analysis methods is whether or not a combination of a specified quantum neural network architecture and the available training algorithms and measurement resources can be optimised. Gradient variance analysis looks at how foundationally the estimators of gradients behave in the space of parameters, finding barren plateau phenomena, where it is observed that the estimators of the gradient exponentially shrink with the size of the system or the number of layers in the circuit itself. Such studies tend to use theoretical methods in random matrix theory and statistical mechanics to make analytical predictions on the gradient scaling of large quantum systems and numerical simulations on particular architectures and problem instances. Classical neural network theory has been modified into local minima and loss landscape analysis techniques to study the geometry of optimization of quantum neural networks, however the stochasticity introduced by measurements and the high dimension complex parameter space poses special problems. Information-theoretic measures of flow Trainability measures defined in terms of quantum Fisher information give characterizations of the sensitivity of quantum circuit responses to parameter changes, and have relations to both gradient scaling and fundamental limits of parameter estimation amid measurement data.

Methods of architecture search inspired by neural architecture search (NAS) in traditional machine learning offer mainstream mechanisms of finding good quantum neural network architectures instead of considering only hand-produced designs. In evolutionary architecture search, genetic algorithms or other evolutionary computation methods are used to search mazes of quantum circuit architecture, using the performance of training tasks or alternative surrogate functions such as repressibility as the fitness. Architecture search This algorithm applies reinforcement learning (RL) methods to architectural design of quantum circuits, viewing quantum circuit design as a sequence of decisions where RL agent policies decide the gates to use, and their qubit targets as

well as the qubit connections within the circuit, based on the gates in the library and design size. Bayesian optimization methods represent the correlation between the architectural parameter and the performance measures based on probabilistic surrogate models that are used to search architectural spaces quickly and trade off between exploration and exploitation. Architecture search methods based on gradients use architectural choices as differentiable parameters (which can be optimised with circuit parameters) but the discrete nature of most architectural choices and quantum gradients being hard to estimate mean that the methods cannot be used directly in quantum systems.

Simulation frameworks Tensor network simulations are the keys to classical analysis and validation of quantum neural network architectures that rapidly prototype and explore more theoretical space without a quantum externalizer. State simulation Matrix product state (MPS) simulations are well suited to quantize representation and manipulation of limited entanglement quantum states, and can be used to scale up to comparatively qubit-scales of quantum circuits, at least to shallow circuit definitions, and classically. Tree tensor network (TTN) and projected entangled pair state (PEPS) methods can be used to simulate other circuit structures and entanglement patterns, and gain computational cost in a variety of ways depending on entanglement structure of the quantum circuit being simulated. Canonical algorithms Based on circuit architecture Canonical algorithms: Tensor network contraction algorithms are optimised in terms of circuit architecture, including slicing algorithms, as well as approximation algorithms which have rigorous error bounds on the quality of approximation. Such simulation techniques do not simply permit the validation and debugging of quantum neural network implementations but also give understanding of the nature of architectural features of quantum advantage as they demonstrate when only classical simulation is feasible.

The techniques in quantum complexity theory can offer critically important tools of analysis into the capabilities and constraints of quantum neural network architectures in computation. Minimum depth bounds on quantum circuit complexity theory can help to determine minimum depth bounds on the implementation of certain quantum gates or the preparation of certain quantum states that specify an architectural design (quantum architecture by determining minimum requirements on irreducible resources) are found using minimum depth bounds. Theoretical frameworks In theories related to the understandability of the representational capabilities accessible under various architectural depth constraints, computational complexity classifications of quantum circuit families, such formed by categories such as constant depth quantum circuits, logarithmic depth circuits and polynomial depth circuits, are considered. Methods to analyse the presence of a provable computational benefit in selected quantum neural network architectures often use the methods of both polynomial hierarchy analysis and

quantum advantage analysis the methods bridge the gap between the property of the architecture and complexities-theoretic separations. These theoretical approaches can be seen to supplement empirical analysis since they put up basic constraints and possibility findings which lead to practical development in architecture.

The noise analysis and error modelling techniques determine the sensitivity to different error mechanisms existing in quantum hardware implementations and their dependence on the architectural decisions. The error analysis of a gate error tests the accumulation of error due to each execution of the circuit and the overall error is generally proportional to the depth of the circuit and dependent on the order in which side gates are executed and what gates the circuit consists of. Coherence time analysis also considers the execution capability of suggested architectures in the coherence constraints of available quantum devices considering circuit depth and gate operations and measurement time. Error propagation Endures (Its studies) determine the responsiveness of errors put in certain locations within the circuit to the ultimate measurement results, and this may indicate architectural design features that prove to be both good and bad at counteracting errors. Channel capacity and fidelity analysis tools generate metrics of the capacity of noisy quantum circuits to transmit or process quantum information in a more information-aware manner, include realistic noise models, as well as can guide architectural optimization of noise robustness.

4.3 Challenges in Quantum Neural Network Architecture Development

Effective quantum neural network architecture model development faces several basic issues that emerge at the interface between quantum physics and computational complexity as well as limited practical hardware requirements. The barren plateau phenomena can be possibly the most extensive issue due to exponentially decaying gradients of randomly initialised parameterized quantum circuits that makes the normal gradient based optimization strategy ineffective with deep / wide quantum neural networks. This has been due to the concentration of measure effects of high-dimensional quantum space, in which random quantum circuits create a state whose associated properties are concentrated near its means, with exponentially small changes that gradients can potentially avoid, in error. The problem of barren plateaus does not just represent a convenient problem with optimization, but it is a core part of expressive quantum circuits, posing a serious tradeoff between expressiveness and representational power, and the trainability needed and learnability desired. The architecture of strategies to counter the barren plateau has been approached by structured initialization schemes adding problem relevant information, layer based training protocols where networks are constructed in layers and architecture limitations such as locality or symmetry which reduce the chances of random behaviour with preserved useful expressibility.

Nonetheless, it has so far been unable to come up with a comprehensive idea of what architectural properties allow both high expressibility and trainability to large problem domains, which is currently an area of research but has large scaling implications to quantum neural networks.

Applications noise and errorhalation Hardware noise and error accumulation are widespread issues and inherent to investigate the invasiveness and sophistication of deployable quantum neural network models on near- and current quantum hardware. The imperfect control pulses that arise as a result of a gate error, cross talk between qubits, and environmental interactions add up but make no additional contributions to the overall error even during circuit execution, and its overall effect is often proportional to the number of gates used in the circuit. Energy relaxation and dephasing are examples of decoherence processes that place a hard time limit on quantum computations and the execution times of quantum circuits are limited by the coherence times between microseconds and milliseconds, depending on the type of qubits used. Measurement noise is an additional noise in the critical point of learning classical computation in quantum computation, and may influence training signals, and may bias the optimization. These sources of errors are complexly interacted with the architectural decisions: deeper circuits contain more error at the gate, and most need longer coherence times, whereas larger circuits can be applied to fewer tasks with shallow implementation, but require more resources and different types of error paths. This problem is made worse by the fact that error rates and error properties can be different in different qubit technologies, quantum processors, and even qubits of the same device, and architectural designs must be both flexible enough to deal with heterogeneous error profiles or be robust enough to operate in a variety of noise conditions.

The architectures of the existing qubit hardware have an important limitation in the connectivity of qubits to each other, particularly when it is thought that all qubits interact with each other, or anything can be put anywhere in the circuit. The majority of quantum processors put qubits in one-dimensional chains or two-dimensional grids or other restricted topologies with only geometrically nearest qubits able to interconnect with a two-qubit gate. The imposition of quantum neural network structures with non-local interactions makes the addition of SWAP gates to transfer quantum information around the processor substrate structure potentially multiply the circuit depth by factors related to the graph distance between qubits of patterns of the desired and available connectivity. There is a basic trade off between such connectivity challenge: an architecture with full connectivity gives the optimum possible expressiveness, but can be compiled into hardware with narrow connectivity, or the architecture can be designed with hardware topology constraints in mind at the expense of expressiveness, or vice versa. It is also made more complex by the fact that various quantum computing platforms adopt dissimilar patterns of connectivity, which means there is no universal architectural

answer, and platforms-specific architectural designs must be developed or an extremely refined compilation process performed to transmute a general architecture into a variety of hardware limitations.

A particularly challenging issue with quantum neural networks to load quantum data is when quantum data is to be fed input in a format of classical information and the process of loading the data into quantum states is incurring as many circuit depths as the follow up quantum processing. Amplitude encoding, providing a compact classical-to-quantum superposition encoding of data, is $O(n)$ -gate amplifiable to n -dimensional data, and can breakdown control and error dominated quantum circuits. This difficulty is not only technical but also faces the underlying question concerning why quantum neural networks can deliver benefits in working with classical data: whether on-loading data into quantum format requires resources as much as classical processing, the quantum computation itself must give significant benefits so as to obtain overall benefit. Various architectural methods have tried to solve this by repeatedly applying the same quantum encoding to many training examples, or by quantum random access memory (QRAM) structures, however QRAM is inherently very difficult to implement and raises serious questions of scalability. Other architectural models such as quantum kernel methods and quantum feature mapping are aimed at minimising the overhead of loading and encoding data by implementing quantum circuits that are tailored to be as simple as possible (minimising encoding complexity), but also limits the kind of quantum processing that can be applied afterward.

Small numbers of qubits of the available quantum hardware tend to pose a fatal degree of limitation to the width of quantum neural network design spaces, and will remain so in general once qubit counts increase since the size of problems and the capacity of model wanted also increase. Applications such as quantum chemistry simulation, large scale optimization and high-resolution image processing can need hundreds to thousands of qubits to be reachable, which are much larger than current NISQ devices. This size constraint is incompatible with a depth constraint created by noise and coherence: in the event that a lack of qubits prevents parallel processing, parallelization causes depth to rise, and methods to limit depth by parallelizing processing force even more qubits. Innovations in architecture that seek to mitigate the ordeal through small width involve quantum circuit cutting where large circuit has been factored into small subcircuits that can be executed on limited hardware, with answers being combined by classical post-processing but this method scales exponentially with the number of cuts. Another method, which may allow meaningful computations with only small numbers of qubits on smaller subsets of large problems by processing information at various granularities, is hierarchical and multi-resolution architectures, which may heed representative and valuable information with bunched quantum representations.

The quantum neural net measurement bottleneck is due to the probabilistic nature of quantum measurement and the fact that, when measuring quantum states, that action collapses superpositions causing the need to re-run the circuits many times in order to estimate expectation values. To achieve a single expectation value value with error ϵ , one would generally need $O(1/\epsilon^2)$ samples and this sampling scale is multiplied by more complicated loss functions involving more than one observable. In quantum neural networks with thousands of parameters, which need to be estimated by parameter-shift rules to compute the gradient, the overall measurement cost may be prohibitive and be the dominant factor in either training time or resource usage. The use of architecture plays an important role in measurement requirements: the cost of circuits that condense the information of interest into an easily measurable observable is lower than the cost of circuitry that generates distributed information that is harder to measure using multiqubit interactions. This measurement issue generates a fundamental asymmetry between quantum and classical neural networks because to compute the output of a classical network, one must make only a single forward pass through the network, whereas to compute the distribution of neural network outputs or expectation values one must communicate with a statistical sample, as is done with quantum networks.

Table 1: Quantum Neural Network Architectures – Applications, Techniques, and Methods

Sr. No.	Application Domain	Techniques Employed	Methods for Implementation	Primary Challenges	Key Opportunities	Performance Metrics
1	Quantum Chemistry Simulation	Variational quantum eigensolvers, unitary coupled-cluster ansatz, symmetry-adapted circuits	Hardware-efficient ansätze, qubit tapering, fermion-to-qubit mapping	Limited qubit count, circuit depth constraints, state preparation complexity	Near-term quantum advantage, industrial relevance, natural quantum structure	Energy accuracy, convergence rate, quantum resource count
2	Molecular Property Prediction	Quantum graph neural networks, quantum feature maps, quantum kernels	Message-passing circuits, atomic orbital encoding, hybrid quantum-classical models	Encoding classical molecular structures, limited training data, noise sensitivity	Drug discovery applications, materials design, property optimization	Prediction accuracy, sample efficiency, inference cost

3	Financial Portfolio Optimization	Quantum approximate optimization, variational circuits, quantum annealing-inspired architectures	Binary encoding of asset allocation, QAOA layers, constraint embedding	Demonstrating advantage over classical optimization, problem encoding overhead	High-value industrial applications, natural fit for optimization problems	Solution quality, time to solution, quantum resource requirements
4	Generative Modeling	Quantum generative adversarial networks, quantum circuit Born machines, quantum autoencoders	Parameterized generator circuits, quantum-classical discriminators, variational training	Training instability, mode collapse, measuring generative quality	Quantum state learning, probability distribution modeling, synthetic data generation	Fidelity to target distribution, sample quality, KL divergence
5	Image Classification	Quantum convolutional networks, quantum feature extraction, hybrid quantum-classical models	Qubit encoding of images, local quantum circuits, pooling via measurements	Classical data encoding overhead, limited resolution, noise accumulation	Computer vision applications, potential quantum advantage with quantum images	Classification accuracy, quantum circuit depth, encoding efficiency
6	Combinatorial Optimization	Graph-based quantum neural networks, quantum alternating operators, structured ansätze	Problem-specific Hamiltonians, constraint-preserving circuits, variational optimization	Problem embedding complexity, local minima in optimization landscape	Logistics, scheduling, constraint satisfaction applications	Approximation ratio, constraint satisfaction rate, computational resources
7	Time Series Prediction	Quantum recurrent networks, quantum long short-term memory,	Sequential quantum processing, quantum memory mechanisms,	Maintaining coherence across timesteps, limited memory	Financial forecasting, signal processing, dynamical system modeling	Prediction accuracy, memory capacity, temporal dependency capture

		temporal quantum circuits	hybrid recurrence	capacity, gradient flow issues		
8	Quantum State Classification	Quantum-to-quantum learning, measurement-based classification, variational state discriminators	Direct quantum input, optimized measurement bases, quantum kernel methods	Limited labeled quantum data, measurement optimization, state tomography overhead	Quantum technology applications, quantum computing verification, quantum sensing	Classification fidelity, measurement efficiency, generalization performance
9	Drug Discovery	Molecular simulation, binding affinity prediction, generative molecular design	Quantum molecular representations, variational property prediction, hybrid models	Molecular size limitations, accuracy requirements, validation challenges	Pharmaceutical applications, de novo drug design, multi-objective optimization	Binding affinity accuracy, molecule generation quality, computational cost
10	Natural Language Processing	Quantum text encoding, quantum transformer architectures, quantum attention mechanisms	Token embedding in quantum states, entanglement-based attention, hybrid processing	Sequence length limitations, encoding complexity, proving quantum advantage	Language understanding, semantic processing, efficient transformers	Language modeling perplexity, task-specific accuracy, quantum resource scaling
11	Quantum Error Correction Decoding	Syndrome processing networks, neural decoders for quantum codes	Syndrome-to-correction mapping, trainable decoding strategies	Real-time decoding requirements, decoder accuracy, scalability	Improved quantum computing fidelity, adaptive decoding, code-agnostic approaches	Decoding accuracy, logical error rate, decoding latency
12	Materials Science Simulation	Electronic structure calculation, phonon spectrum prediction,	DFT surrogates, quantum embedding methods,	System size limitations, accuracy for complex	Materials discovery, accelerated simulation workflows	Property prediction accuracy, computational speedup, materials

		phase transition modeling	multi-scale modeling	materials, computational cost		space coverage
13	Quantum Sensing Data Analysis	Sensor output processing, adaptive sensing protocols	Quantum neural processing of sensor data, measurement optimization	Real-time processing, sensor noise, limited quantum memory	Enhanced sensing precision, adaptive protocols, multi-sensor fusion	Sensing precision, estimation variance, protocol efficiency
14	Reinforcement Learning	Quantum policy networks, quantum value functions, quantum actor-critic architectures	Quantum state encoding, variational policy optimization, hybrid RL algorithms	Exploration challenges, credit assignment, state encoding complexity	Robotics, game playing, quantum control, efficient exploration	Cumulative reward, sample efficiency, policy quality
15	Anomaly Detection	Quantum autoencoders, quantum density estimation, reconstruction-based detection	Quantum feature learning, variational compression, quantum distance-based outlier detection	Defining normality, threshold selection, computational cost	Cybersecurity, fraud detection, system monitoring, quantum diagnostics	Detection accuracy, false positive rate, computational efficiency
16	Quantum Communication Protocol Optimization	Protocol parameter optimization, routing, entanglement distribution	Variational protocol learning, quantum network simulations	Protocol complexity, fidelity requirements, scalability	Quantum internet development, adaptive communication protocols	Communication fidelity, protocol efficiency, entanglement rate
17	Climate Modeling	Weather pattern recognition, climate prediction, extreme event forecasting	High-dimensional data encoding, temporal modeling, hybrid simulation	Large data volume, long prediction horizons, validation issues	Improved climate prediction, policy applications	Prediction accuracy, forecast horizon, computational resources
18	Recommendation Systems	User preference modeling, collaborative filtering	Quantum collaborative filtering, quantum matrix factorization	Scalability, cold-start problems, data encoding	Personalized systems, efficient similarity computation	Recommendation accuracy, coverage, computational efficiency

			, hybrid recommenda tion			
19	Quantum Circuit Compilation	Circuit optimization, gate synthesis, topology mapping	Neural compilation networks, reinforceme nt learning- based compilation	Optimality guarantees, generalizat ion, computatio nal cost	Automated quantum software developme nt, efficient circuits	Compiled circuit depth, gate count, fidelity
20	Bioinformati cs	Protein structure prediction, genomic sequence analysis, systems biology modeling	Quantum sequence modeling, structure prediction networks, biological network analysis	Data encoding complexity , biological scale, validation challenges	Personalize d medicine, biological discovery, drug targets	Prediction accuracy, computational cost, biological validity

Scalability and generalisation are related issues because quantum neural network representations will not be as small as proof-of-concept demonstrations as systems that aim at solving practical problems on large quantum computers [30-32]. Theoretical insights into size-dependent scaling of the properties of architectural systems are still not exhaustive: although some studies indicate that scaling behaviours grow exponentially with the number of qubits, or architectural structures, with size, others have hinted at the possibility of a polynomial, or even even logarithmic, scaling of the behaviours of certain problem categories with system size. Most quantum neural network demonstrations are done empirically with tens of qubits and shallow circuits, thus the question of whether successful small-scale architecture styles will scale or give way to new qualitatively different methods are open. High dimensionality The high-dimensionality of the complex parameter spaces, the entanglement influence on the representational capacity of quantum neural networks, and the interactions between quantum measurement, classical optimization and learning are all complicating, in general, the generalization analysis, i.e. the behavior of quantum neural networks when trained with finite data. The classical learning theory is not necessarily directly applicable to quantum systems and the creation of quantum specific generalisation bounds that consider architecture properties is still being a theoretical challenge.

The classical simulation efficiency paradox poses the minor but significant difficulty on quantum neural network architecture construction: as long as a quantum neural network architecture can be efficiently simulated on a classical computer, it cannot probably offer quantum advantage, but classical simulation is required to develop, debug and test architectural concepts. Such a paradox has a practical implementation in that researchers

are torn between architectures which are easy to analyse classically but which may be limited in quantum advantage, against architectures that are hard to simulate but understandable only phenomenologically or empirically. The nearest solution to this, provided by tensor network loops, is the possibility to easily perform simulation of quantum circuits with a given structure of entanglement, but the most promising quantum advantage-scale architectures (those with long-range and highly entangled structure) are the most resistant to causal simulation. This has statistical impacts methodologically on how architectural studies are prejudiced, so that they are restricted to designs which can be characterised classically but which may not offer the most encouraging prospects of quantum advantage.

The difficulties between quantum system and classical system component during integration pose an architectural problem in hybrid schemes of combining quantum and classical processing. Latency costs are incurred in the quantum-classical interface because quantum measurement outputs should be sent to classical processor units, which may be computed upon and feed back as updated parameters to be further used in running quantum circuits. This latency has an impact on the feasibility of closely-interlaced quantum-classical architectures as well as the granularity of quantum and classical computing, which can be interleaved. The conceptual issues and practical drawbacks of coherent hybrid architecture designs are different computational models: discrete quantum gates and continuous classical operations, probabilistic and deterministic quantum measurements, and probabilistic and deterministic classical computations. Another challenge in integration is balancing quantum and classical resource allocation quantum resources are valuable and are prone to error, implying that quantum resources should be minimised, whereas quantum advantage may necessitate significant quantum computation, implying that quantum resource utilisation should be maximised. The choices the architect suggests in this regard have a significant implication on the achievability of quantum benefit as well as the feasibility of quantum neural network systems in practise.

The absence of standardisation among the quantum neural network research community poses a problem of comparing architectural implementation methods, remaking the results, and extending existing work. The definitions of some of the important concepts such as depth, expressibility, and quantum advantage are used differently by different researchers meaning that literatures synthesis and meta-analysis may be difficult. Benchmark tasks are not yet standardised and various studies use varied datasets, formulations of problems and performance measures and hardware platforms thus rendering it challenging to compare them well. The development of quantum neural network software frameworks and tools is changing fast and there are numerous competing platforms that offer various abstractions, gate sets and compilation methods and this is likely to result in implementation-based behaviour that masks underlying

architectural concepts. The problem of standardisation is also complicated because the ability of quantum hardware and algorithm implementation methods rapidly advances and can quickly make certain benchmarks or practise implementation redundant. Agreeing on methodologies of architectural evaluation, benchmark suites, and reporting standards, at the community level is one priority of the field growing up and to allow systematic architectural development.

4.4 Opportunities in Quantum Neural Network Architecture Research

The fast-changing world of quantum neural networks architecture offers a great potential of an unprecedented fundamental progress and practical use of quantum hardware capacity and theoretical knowledge increasing. The closest application benefit is near-term quantum advantage, which is the claim that quantum neural networks can deliver useful benefits to specific tasks, notwithstanding noisy intermediate-scale quantum hardware. A variety of applications of quantum chemistry, in general variational quantum eigensolvers to solve the molecular ground state and the dynamics of a quantum system, show access to near-term benefit as quantum computational resources increase to 50-100 qubits with reduced error rates. It has an opportunity in discovering and optimising quantum neural network designs especially in these advantage-ready models; on problem instances in which quantum effects in the target system match quantum circuit constraints and problem instances where classical algorithms have fundamental rather than engineering challenges. The various architectures that are designed in collaboration with domain-related application requirements with domain knowledge with structured ansatzes, symmetry constraints, and problem-sensitive measurements provide the opportunity to achieve the maximum performance within the existing hardware limits as well as provide proof points that encourage further development of quantum hardware.

The opportunities of the hybrid quantum-classical architectures in terms of using the complementary advantages of quantum and classical computing platforms and eliminating the weaknesses of each platform are fascinating. Hybrid architecture design space is mostly uncharted, and there are possibilities to invent new integration patterns that are not based on existing design paradigms such as quantum kernels or variational circuits with classical optimizers. There are possibilities of quantum-classical co-design in which architectural choices are optimally joint to both quantum circuit design and classical neural network design to pre-process quantum inputs, post-process quantum outputs, or create components of other complex computational programmes. They can be combined with hierarchical hybrid designs that use quantum processing to specialise computational bottlenecks within larger classical systems, leaving another promising direction that could use quantum devices to realise quantum advantage on larger tasks

even where a smaller subproblem needs to be solved using quantum devices. The optimum separation of quantum and classical should also change as quantum hardware capabilities increase, enabling the use of adaptive hybrid architectures, which can effectively adapt to both available quantum resources and the features of the problem they are being used.

Research in the development of quantum neural network architectures that are resistant to noise has fundamental opportunities since quantum error correction is too far to be practically used, and noise adaptability is required and not optional. It has opportunities to architecturally construct architectures that are resistant to the dominating error mechanisms of a particular quantum hardware implementation, by using architectural-level symmetries, redundancy, or error-correcting codes as opposed to using architectural-level error correction layers. Suppression of variational error suppression methods with noise mitigation integrated into the trainable quantum circuit architecture are a promising direction to learn robust representations with noise, which utilises parameterization to counter known error channels or mitigate noise. Noise-resilience architecture search, instead of noiseless performance, provides possibilities in finding circuit structures that are resilient to hardware imperfections, which may in turn provide design ideas in the NISQ era that are not the same as the ones applicable to fault-tolerant quantum computers. The interaction between noise and entanglement enables architectures that selectively grow entanglement in order to strike a balance between expressibility and noise accretion, and maximise the trade-off instead of expressibility, insensitive of sensitivity to errors.

The quantum neural architecture search is a frontier possibility to explore effective structures of quantum circuits by automation, and potentially even in most cases, discover non-intuitive structures that can be more effective than those defined with hand-created neural architecture. The discrete search spaces of quantum circuit architectures are high-dimensional, and they are expensive to evaluate on quantum hardware, which opens the possibility of designing highly specialised architecture search methods, which pose structural and environmental quantum-dependent structure and constraints. There are prospects of exploiting classical simulation to search the space of architectures quickly, to run all applicants through simulation prior to quantum hardware testing, but again this strategy needs to take into consideration the fact that, within the definition of quantum advantage, classically-simulable architectures can be found. The Pareto-optimal architectures that are the best trade-offs of various hardware platforms and applications could be found using multi-objective architecture search that considers and optimises both performance, circuit depth, number of gates and noise resilience. Transfer learning The theme of transfer learning in which architectures learned on similar tasks or small scale problem instances drive the search with new tasks allows opportunities to distribute the cost of expensive architecture search across multiple application.

Theoretical principles of learning and interpretation of quantum neural networks offer significant possibilities both to generic quantum information science and application to generic quantum machine learning. There are prospects of thoroughly working out complete architectural theory linking the architectural character of circuits depth, width, entanglement structure, gate composition to learning theoretic measures such as sample complexity, generalisation limits and computational complexity. Machine learning theory Quantum advantage theory is still incomplete, with the possibility of building rigorous learners separations between quantum and classical learning capabilities to larger problem classes than those currently understood, and may lead to understanding application areas where quantum methods can have provable advantages. Knowledge of the correlation between quantum advantage and architecture would be precious informational input in the design of architecture by considering the development work to be directed towards architectural patterns that are most apt to provide advantage. As well as opportunities to mutually refine ideas and techniques which can simultaneously develop all these areas, quantum neural networks have connexions with other fields of quantum information such as quantum error correction, quantum complexity theory and quantum Shannon theory.

Expansion of application domain is an opportunity to recognise new problem domains in which quantum neural networks architectures may add special value beyond the present areas of interest. The application of quantum sensing and quantum metrology Systems based on quantum neural networks processing quantum measurement data to isolate classical information directions are some promising fields where quantum processing is inherently needed. The optimization of protocol used in quantum communication and cryptography systems may utilise quantum neural networks to optimise the protocol, eavesdropping detection, or characterise the quantum channel. There are applications in scientific computing outside the field of quantum chemistry, such as materials design, drug discovery, high-energy physics simulation, and cosmological modelling, where quantum effects are critical to the underlying research question, and quantum neural networks can be used. New classical machine learning systems such as foundation models, multi-modal learning and continual learning generate the prospect to examine whether quantum strategies can resolve issues such as catastrophic forgetting, and few-shot adaptation or extreme level of computational efficiency that classical systems fail to achieve.

The opportunities offered by standardisation and benchmarking initiatives include boosting quantum neural network research by coordinating communities and having a shared infrastructure. There is an opportunity to develop a standardised benchmark suites covering different types of tasks, the sizes of problems, and metrics of evaluation in order to adequately compare the architectural methods and obtain the developmental progress over time. Open-source software frameworks that enact architectural building

blocks, training algorithms, and evaluation tools could make quantum neural network work more democratic and minimise the amount of overlapping work done in research groups. Provision of shared programmes, and cloud quantum computing platforms, provide access to shared quantum hardware by researchers who do not have access to specialised quantum hardware, but standardisation of evaluation programmes across heterogeneous hardware platforms is problematic. The establishment of community norms regarding reporting, the ability to replicate work, and open science practises would help hasten the process by making research reliable on previous research instead of relying on already known outcomes or working on paths that have already proven unproductive.

The co-design of quantum hardware opens the possibility of matching quantum neural network architecture design with the design of quantum computing hardware, and provides feedback between hardware capabilities and requirement of algorithms. As quantum hardware designers are taking decisions on qubit connexions, gate sets, and control circuits and measurement system design, quantum neural network researchers might consult on the best architectural characteristics to enable hardware to be made. On the other hand, new quantum hardware options such as mid circuit measurement, qubit reuse and feed-forward control open up possibilities to new quantum neural network designs that can make use of these capabilities in ways that were not possible on previous generations of hardware. Quantum machine learning-focused application-specific quantum processors can be considered a longer-term opportunity, which might be more efficient to run quantum neural networks compared to general-purpose quantum computers, but that aspect of such specialisation necessitates enough maturity in identifying which architectural patterns are universally applicable when running quantum machine learning code.

Quantum neural networks hold education and workforce development opportunities to develop human capital and create the interdisciplinary knowledge needed to carry on with the development. There are possibilities to build educational opportunities, courses, and training programmes that can effectively merge quantum physics, machine learning, and practical programming skills and equip the scientist to be involved in the creation of quantum neural networks. Competitions, hackathons and challenges based on quantum machine learning may bring in larger groups of people to the field, which may identify well-designed researchers with unconventional backgrounds and find new ways of solving longstanding issues. Forming communities of practise to incorporate quantum physicists, machine learning experts, domain scientists, and hardware engineers provides a chance to cross-pollinate ideas and then collaborate in solving problems which no one field alone could do. With the continued development of quantum neural networks into a useful technology, there will be prospects of applied training programmes that train

practitioners to use quantum machine learning tools in practise in industry, government, and science.

4.5 Impact of Quantum Neural Network Architecture Innovations

The effects of breakthroughs in the field of quantum neural network architectures can be traced along various levels, including the level of underlying science to the applied technological functionality and implications on the society in the general. Research On quantum computing research itself, quantum neural network architectures have sparked novel views of quantum circuit design, variational quantum algorithms, and noisy quantum device capabilities in the near future. The knowledge obtained through studying the problem of trainability such as that of barren plateaus has provided deeper insight into the field of high-dimensional quantum state spaces and concentration phenomena which has made contributions to quantum information theory beyond the use of quantum machines in machine learning. Architectural inventions such as symmetry-adapted ansatzes, entanglement structuring techniques and measurement optimization have had broader impact on quantum algorithm design with concepts developed to quantum neural networks applied in quantum optimization, quantum simulation and other fields of quantum algorithm design. Empirical value These empirical implementations have found information about which hardware is necessary and where is a bottleneck, which has shaped quantum computing roadmap and hardware development priorities at quantum computing giants companies and research centres.

The diversity of influence on machine learning theory and practise is reflected in the new conceptualization, algorithm strategies, and revelations regarding learning and representation. The quantum neural networks have led to rethinking about what a neural network is and the classic understanding of the concept of a neuron and connexion has been extended to include quantum mechanical information processing. Theoretical studies of quantum learning have shown new complexity separations, and provided evidence of learning problems in which quantum learning strategies outperform classical ones, which is part of the computational learning theory and the understanding of learning fundamental limits. The practical effects are still less extensive since hardware is a limitation, however, demonstrations in fields such as quantum chemistry and generative modelling provide evidence of possible applications that can become groundbreaking with the growth of quantum hardware. There have been positive outcomes in the quantum / classical machine learning interface with quantum interventions inspiring new classical algorithms (quantum-inspired optimization, tensor network machine learning) and classical algorithms (batch normalisation, attention mechanisms) being tailored to the quantum context.

Fields where quantum neural network architectures have been applied to gain a scientific impact include especially quantum chemistry and materials science, fields where quantum system simulation is transforming computational methods. VQE which uses quantum neural network architectures has been used to compute ground state energies of larger and larger molecular systems in the protocol as close as it gets to that of classical methods, and even surpass them in the accuracy of strongly correlated systems. Its effects are not limited to benchmark computations, but also have an overall effect on how chemists solve computational problems, and quantum-classical hybrid workflows are nowadays considered a staple of computational chemistry research. Quantum neural network-based property prediction and molecular optimization components are starting to be included in both drug discovery and materials design processes, but both efficiency and robustness in bigger quantum computers are needed before quantum computing can have a large-scale impact. These applications and how they were made possible by the architectural innovations that have been described such as chemistry-aware ansatzes, noise reduction methods, hybrid quantum-classical optimization are constitutive of methodological contributions to the future health of computational chemistry practise within the scope of the increasing abilities of the quantum hardware.

Quantum neural network technologies are starting to have a premature effect in the industrial and commercial fields, mostly on the level of exploratory research and development and not on deployed production-capable systems. QNNs have been studied by financial services companies to optimise portfolios, examine risk, price derivatives, and more, and several large companies have formed groups of quantum computing researchers and collaborated with quantum hardware manufacturers. The effect up to this point is mostly preparatory, to establish expertise, experiment with use cases, prove the potential benefits, but this base puts those that adopt it early in a position to utilise the capabilities of quantum as they come to fruition. The pharmaceutical and chemical sectors have been funding quantum neural network applications on drug discovery and materials design problems due to the inherent quantum aspect of the systems being studied and the capacity of the classical simulation to simulate them. Another area of impact is applications in logistics, manufacturing, and operations research that make use of quantum neural networks to perform combinatorial optimization, although it has been seen that it can make impressive improvements on many problem instances in comparison with highly optimised classical algorithms.

Exposure to quantum neural network research has educational effects such as training, educational and interdisciplinary curricula bridging between quantum physics and machine learning. Courses and degree programmes in quantum computing and quantum information science have been created in universities around the world and quantum machine learning and quantum neural networks are often mentioned as examples of how quantum computing can be used. This pedagogical effect requires not just the role played

in formal education, but in a wider scope of scientific literacy, where quantum neural networks can be used as easily reached points of entry into learning the potential use of quantum computing. Online courses and a series of in-person courses in quantum programming have involved thousands of students and employees to create a labour force with appropriate skills and increase the number of individuals with whom quantum neural network development can take place. It has had especially profound effects on interdisciplinary training of researchers in the field, since quantum neural networks necessarily involve the presence of expertise in quantum physics, computer science, mathematics, and most application domains demand the use of domain-specific knowledge, which encourages educational methods that dissolve barriers to discipline.

More general effects on the commercialization of quantum computing and quantum industry ecosystems are the contribution of quantum neural networks to the demonstration of near-term quantum computing applications and spurring further investment. One of the most commonly mentioned areas of business case applications of quantum computing is quantum machine learning that has gained popularity with quantum neural networks, impacting funding decisions and overall strategic look at quantum computing companies. Quantum neural network applications have contributed to the development of quantum machine learning software platforms, cloud-based quantum computing services, and quantum programming frameworks and provide infrastructure that is useful to the rest of the quantum computing community. Quantum neural network applications are frequently used as demonstrators and early use cases in partnerships between quantum hardware vendors, cloud computing vendors and the end user industry, helping to transfer technology and build the ecosystems. This has had a both positive and negative effect on the timelines and expectations of quantum computing, quantum neural networks have both fueled excitement about realistic uses in the near term and helped improve clarity regarding the needs, such as noise sensitivity, and difficulty due to challenges in establishing unambiguous quantum advantage.

There are methodological implications both to the practise of computational scientific research, such as quantum-classical hybrid workflows introduced in quantum neural network research having influence on the wider understanding of computational paradigms. Bringing quantum computing resources to machine learning circuits and using quantum computers as specialised accelerators of particular computational tasks and coming up with interfaces between quantum systems and classical systems are methodological innovations that are not restricted to quantum neural networks. The noise-conscious algorithm design philosophy, where the realisation of the hardware's imperfections is considered during its design and not as an implementation fact, has manifested a methodological shift that can be applied to quantum computing. Quantum neural network practises of reproducibility and release (of quantum circuit implementations), attentiveness to characterising hardware platforms, and statistical

analysis of measurement data, are part of wider debates on computational reproducibility and open science.

Although to a large extent still in the future, since the technology is in its initial stages, societal and ethical considerations of its quantum neural network development are starting to be taken into consideration. Issues related to the possible attack on cryptographic security by quantum computing have been raised to the use of quantum machine learning, involving the possibility of quantum neural networks being more efficient at cryptanalysis or the possibility of quantum-safe machine learning systems. The questions of access and equity, who is going to use quantum computing strengths and whether quantum technologies will lead to more or less significant technological disparities will be relevant to quantum neural networks as they are relevant to quantum computing in general. The considerations of the environmental aspects of the energy usage and cooling needs of quantum computers that perform quantum neural networks are a part of a wider debate concerning the sustainability of computation technologies. The fact that quantum neural networks could be used to solve significant scientific and societal problems, including drug discovery or climate simulations, is a potential of solutions that indicate positive impact and drive further progress while bringing the ethical question of priorities and resource distribution to the development of quantum technology.

4.6 Future Directions in Quantum Neural Network Architecture Research

The interaction between the development of better quantum hardware, increased theoretical knowledge and new problem domains dictate the future trends of the study of quantum neural network architecture and provides a rich environment full of research opportunities and with the potential to make a breakthrough. Fault-tolerant quantum neural networks imply one of the key areas of future technology that will be even more of a significant issue once quantum error correction techniques become fine-tuned enough to allow architectural designs that are fundamentally different than those accessible today. Architectural limitations and possibilities The move to noisy to error-corrected quantum computation will fundamentally change the constraints and opportunities associated with depth, which will be reduced significantly. It will be possible to explore far deeper quantum neural network heavily relying on portability The representational capability of a much deeper quantum neural network can far outstretch that of shallow NISQ circuits. But quantum error correction will impose new overhead costs, which means that quantum error correction will impose resource constraints on a system already constrained by its physical resources, and required to be overcome by new architectural designs that can use quantum error-corrected qubits effectively and may offload suitable computational tasks to less-guaranteed resources or even classical

resources. Future research directions in fault-tolerant quantum neural networks will involve understanding the best architectures to use when making error-corrected quantum computers, trying to develop compilation and optimization methods that consider the cost of error correction, and trying to find out whether those principles used in architecture development in the NISQ age apply to the fault-tolerant world.

QNNs and large-scale learning Foundational models QNNs approach A proposed future development direction inspired by the disruptive nature of large language models and other foundation models in standard machine learning. Questions of interest are whether quantum methods can be useful in training or inferencing billion or trillion-parameter networks, whether quantum-classical hybrids may be helpful in more effective generation of foundations models or whether quantum neural networks can lead to entirely new large-scale pre-training and few-shot learning. These are daunting tasks that will eventually demand extraordinary improvements in numbers of qubits, quality and connectivity, well beyond the present state of the art, in order to realize the full potential, nevertheless the payoff is worthy of the investigation into scalable quantum neural network implementations at the scale of massive neural network models. Future research is more focused on quantum analog instruments to analogous classical architectures, whether quantum entanglement can effectively model the long-range interactions as seen in large complex models, and how training algorithms can be developed to interact with many billions of parameters of quantum neural networks that are trainable despite barren plateaus and other adverse optimization problems.

Table 2: Quantum Neural Network Architecture Parameters - Impact and Future Directions

Sr . No.	Architectu ral Aspect	Impact on Performan ce	Implement ation Challenges	Current Best Practices	Emerging Trends	Future Research Directions
1	Circuit Depth	Determines expressibili ty; deeper circuits represent more complex functions	Increased error accumulatio n, longer coherence time requirement s, barren plateaus	Shallow to moderate depth (10-100 layers) matching hardware capabilities	Adaptive depth, layer-wise training, depth-width trade-offs	Theoretical depth requirements, optimal depth for applications , fault-tolerant deep networks
2	Network Width	Affects system size representab	Limited qubit availability,	Width matching problem size	Hierarchical width reduction,	Scaling laws for width,

		le, parallelism available, information capacity	increased resource requirements, connectivity constraints	within hardware constraints	multi-resolution architectures	width efficiency metrics, optimal width allocation
3	Entanglement Structure	Determines correlation capabilities, influences quantum advantage potential, affects trainability	Complex entanglement patterns require deep circuits, hardware connectivity constraints	Structured entanglement (brick-layer, tree-tensor) balancing expressibility and implementability	Adaptive entanglement, learned entanglement patterns, entanglement-aware design	Entanglement-performance relationships, optimal entanglement for tasks, entanglement as resource
4	Gate Selection	Affects hardware mapping, expressibility of gate set, compilation complexity	Limited gate sets on hardware, decomposition overhead, gate fidelity variations	Hardware-native gates prioritized, universal gate sets for flexibility	Continuously parameterized gates, learned gate sequences, adaptive gate selection	Minimal gate sets for universality, gate set optimization, novel quantum gates
5	Parameterization Strategy	Influences optimization landscape, affects trainability, determines model flexibility	High-dimensional parameter spaces, initialization sensitivity, barren plateaus	Structured parameterization, initialization incorporating problem knowledge	Layer-wise parameterization, hierarchical parameters, learned parameterization schemes	Optimal parameterization for tasks, parameter sharing strategies, meta-learned parameterization
6	Measurement Strategy	Determines information extraction efficiency, affects training	Shot noise, measurement basis optimization, multi-	Pauli grouping for measurement efficiency, adaptive	Classical shadows, randomized measurements,	Optimal measurement for objectives, measurement-aware

		signal quality	observable estimation	measurement protocols	measurement learning	architecture design, quantum advantage in measurement
7	Hybrid Architecture Design	Balances quantum and classical processing, determines quantum resource utilization	Interface latency, optimal task division, heterogeneous optimization	Quantum feature extraction or kernels with classical downstream processing	Deeply integrated hybrid models, quantum-classical co-design, adaptive resource allocation	Principled hybrid design frameworks, quantum-classical interface optimization, hybrid advantage theory
8	Noise Mitigation Integration	Affects effective circuit fidelity, enables deeper useful circuits, impacts computational cost	Error mitigation overhead, applicable mitigation techniques, architecture constraints	Error extrapolation, dynamical decoupling, probabilistic error cancellation	Architecture-aware mitigation, learned error mitigation, mitigation-expressibility trade-offs	Native noise-resilient architectures, mitigation integrated in design, fault-tolerance pathways
9	Symmetry Incorporation	Reduces effective parameter space, improves sample efficiency, encodes problem structure	Identifying relevant symmetries, implementing symmetry constraints in circuits	Symmetry-adapted ansatzes for quantum chemistry, conservation law enforcement	Learned symmetries, approximate symmetries, symmetry breaking when beneficial	Automatic symmetry detection, symmetry-expressibility balance, novel symmetry applications
10	Initialization Methods	Affects optimization trajectory, influences	Problem-specific initialization requires domain	Near-identity initialization for shallow circuits,	Learned initialization, meta-initialization from related	Theoretical initialization guarantees, task-

		barren plateau susceptibility, impacts convergence	knowledge, identity initialization limits expressibility	problem-informed initialization when available	tasks, initialization search	agnostic good initializations, initialization-architecture interaction
11	Training Algorithm Choice	Determines convergence properties, computational cost, hardware requirements	Gradient vanishing, high measurement cost, classical optimization scalability	Parameter-shift gradient methods, gradient-free optimization for difficult landscapes	Natural gradients, momentum-based methods, adaptive learning rates	Quantum-native optimization algorithms, improved gradient estimation, meta-learning optimizers
12	Regularization Techniques	Controls overfitting, improves generalization, influences architecture simplicity	Quantum analog design challenges, balancing regularization and expressibility	Penalty on circuit complexity metrics, early stopping, parameter magnitude constraints	Quantum dropout analogs, noise as regularization, learned regularization schemes	Quantum-specific regularization on theory, generalization guarantees, adaptive regularization
13	Connectivity Patterns	Determines implementable gate placements, affects compilation overhead, influences entanglement capability	Hardware topology constraints, non-local interactions require SWAP overhead	Respect hardware connectivity when possible, topology-aware architecture design	Learned connectivity patterns, adaptive routing, connectivity as optimizable parameter	Optimal connectivity for problems, hardware-agnostic connectivity design, connectivity-depth trade-offs
14	Data Encoding Methods	Affects quantum advantage potential,	Encoding complexity can dominate	Amplitude encoding for compact representation	Trainable encodings, problem-specific	Theoretical encoding requirements for

		determines encoding overhead, influences downstream performance	computation, fidelity of encoding, classical-quantum interface	on when feasible, angle encoding for simplicity	encoding design, encoding-processing co-design	advantage, minimal encodings, quantum data structures
15	Architecture Modularity	Enables composition and reuse, affects interpretability, influences development workflow	Module interface design, composition semantics, optimization across modules	Hierarchical architectures, modular variational circuits, quantum-classical modules	Learned module composition, automatic modularization, adaptive module structures	Module libraries, compositional theory for quantum circuits, transfer learning across modules
16	Scalability Mechanisms	Determines behavior as problem size grows, affects path to larger quantum computers	Theoretical scaling understanding incomplete, empirical validation at small scales	Problem decomposition, hierarchical processing, focusing quantum resources on bottlenecks	Automatic scaling strategies, scale-dependent architecture adaptation	Scaling laws for quantum neural networks, provable scalability guarantees, asymptotic behavior theory
17	Expressibility Metrics	Quantify representational capacity, guide architecture selection, theoretical characterization	Multiple non-equivalent metrics exist, relationship to task performance unclear	State coverage measures, entanglement metrics, gate set universality	Learned expressibility predictors, task-specific expressibility, multi-metric characterization	Unified expressibility theory, expressibility-performance connections, minimal expressibility requirements

18	Fault Tolerance Integration	Enables error-corrected quantum neural networks, dramatically changes design constraints	Error correction overhead very high, limited availability of fault-tolerant hardware	Not yet practical for near-term devices, theoretical investigations ongoing	Early explorations of error-corrected quantum machine learning, logical qubit architectures	Architecture principles for fault-tolerant regime, optimal error correction for quantum ML, hybrid corrected-uncorrected designs
19	Attention Mechanisms	Enables selective information processing, potentially improves efficiency, parallels classical transformers	Quantum implementation complexity, measurement requirements, demonstrating advantage	Early-stage research with proof-of-concept quantum attention circuits	Quantum self-attention, cross-attention between quantum-classical, learned attention patterns	Quantum attention theory, efficiency advantages, applications in quantum sequence processing
20	Dynamic Architectures	Adapt structure during execution or training, potentially improve efficiency and performance	Control complexity, maintaining coherence during adaptation, hardware support requirements	Early explorations with conditional circuits, measurement-based adaptation	Learned architecture policies, runtime architecture optimization, data-dependent structures	Theory of dynamic quantum circuits, optimal adaptation strategies, hardware support for dynamics
21	Quantum Memory Integration	Enables complex temporal processing, maintains quantum information	Quantum memory fidelity limitations, coherence time constraints,	Quantum registers serving as short-term memory, measurement-based	Quantum random access memory integration, learned memory management	Optimal memory architectures, memory-computation trade-offs, scaling

		n across operations	memory capacity	memory reset	, hybrid quantum-classical memory	memory capacity
22	Benchmarking Frameworks	Enable fair comparison, track progress, identify promising approaches	Defining meaningful benchmarks, hardware heterogeneity, rapid evolution of field	Task-specific benchmarks, standardized datasets when available, reporting best practices	Comprehensive benchmark suites, normalized performance metrics, open leaderboards	Quantum-advantage-sensitive benchmarks, standardization efforts, benchmark theory
23	Multi-Task Architectures	Enable learning multiple related tasks, improve sample efficiency, enable transfer learning	Task interference, architecture capacity allocation, optimization challenges	Shared quantum layers with task-specific readouts, multi-objective training	Learned task relationships, dynamic task prioritization, continual learning architectures	Multi-task learning theory for quantum, optimal architecture sharing, catastrophic forgetting mitigation
24	Interpretability Features	Enable understanding of quantum network decisions, debugging, trust building	Quantum state complexity, measurement collapse, classical interpretation challenges	Visualization of quantum states, analysis of learned parameters, ablation studies	Quantum attention visualization, learned interpretable representations, causal analysis	Quantum interpretability theory, faithful explanation methods, interpretability-performance trade-offs
25	Hardware-Software Co-Design	Optimizes quantum architectures for specific hardware, enables better	Requires tight integration across stack, hardware diversity, generalization challenges	Hardware-efficient ansatzes, platform-specific compilation, noise-adapted designs	Co-optimization of architecture and hardware, learned hardware models,	Automated co-design frameworks, hardware-agnostic principles that specialize well, co-

performanc e	adaptive compilation	evolution of hardware- algorithms
-----------------	-------------------------	---

Neuromorphic quantum computing is a possible and hypothetical direction in the future with architecture designs that better resemble biologically realistic neural systems but make use of quantum effects. In contrast to the existing gate-based quantum neural networks that use digital quantum circuits, analogue quantum systems or continuous-variable quantum computing or quantum annealing platforms could be used by neuromorphic quantum systems to execute neural network computations more directly. Future directions of research incorporating quantum dynamics Incorporating quantum superposition and entanglements quantum quantum neural networks Incorporating quantum superposition and entanglements quantum quantum neural networks and quantum quantum neural networks can potentially have a computational or efficiency benefit compared to completely quantum neural networks. A different neuromorphic-inspired direction with possible near-term implementation is quantum reservoir computing in which temporal input sequences are processed only by readout layers trained, and fixed random quantum systems made process them. These difficulties have been encountered in being able to map significant neural processes onto accessible quantum hardware platforms and proving to be more advantageous over classical neuromorphic methods that are in turn more economical than conventional digital computing processes at particular tasks.

5. Conclusion

Analysis of neural network architecture dimensions depth, width, expressive power shows a field at the cross ways of theoretical exaltation and practical accomplishment. As this full exploration has shown, quantum neural networks are indeed much more than just quantum equivalents to classical neural networks; they are entirely new modes of computation which can use quantum events to potentially break the classical constraints on computations of certain classes of problem. Architectural decisions that govern those systems, including the micro-level of quantum gates and parameterization schemes or the macro-level of hybrid quantum-classical implementations of workflows, and even the application-specific structure, highly affect the theoretical potential as well as practical execution of the quantum neural network implementation.

The overview of the current research synthesis of this chapter points to both the achievements and the lack of solutions to the problem. Theoretical proposals have been replaced on the progress side by experimental demonstrations on quantum hardware, by foundational architectural patterns such as variational quantum circuits and hardware-

efficient ansatzes, by developed techniques to deal with trainability and noise, and by suggested uses of quantum methods that may be fruitful. Theory Technical theory has elucidated provider-customer linkages amidst the system attributes in a circuit, whereas practise has shown that near-mix quantum devices are or are not very capable of carrying out machine-learning tasks correctly. The overlapping of the quantum computing and machine learning circles has resulted in flourishing interdisciplinary research that has brought forth novel practises that have enhanced both fields.

Nevertheless, quantum neural networks have an uphill road towards the implementation of their transformational promise. The barren plateau phenomenon remains the limiting factor to the expressiveness and trainability of quantum circuits, necessitating architectural designs to be more expressive but achieve optimization feasibility. Quantum hardware such are effects on the number of qubits, error rates, connectivity, and coherence time place strong limits on architectures that can be run, with existing systems able to only run relatively limited-depth circuits with relatively few number of qubits. The problem of quantum advantage, determining the problems in which quantum neural networks are proven to work better than classical models than their counterparts, is still to be solved, and some applications of quantum chemistry are clearly benefited but most applications of classical data processing are still open. Any form of scalability, whether of quantum hardware, or of the architectural principles underlying the scaling which will work as the system grows, is a basic uncertainty which impacts upon long-run prospects.

Chapter 5: Resource Scaling in Quantum Kernel Estimation and Measurement Protocols

1 Abstract

Quantum kernel estimation and quantum measurement protocols Resource scaling, at the interface between quantum information theory, statistical learning and experimental quantum engineering. This chapter analyses the topography of resource demands the number of qubits, circuit level, number of state preparations and size of measurements, complexity of classical post-processing inputs and demands of error avoidance to measure quantum kernels with controlled precision on near-term and heritage quantum platforms. We integrate the base theoretical findings, current randomized-measurement theory like classical shadows, and more recent developments in entangled and adaptive measurement designs and empirical research of concentration and sample complexity with significant implications on quantum machine learning. Dominating algorithmic primitives are summarized in two large tables encompassing their scaling of resource usage, implementation caveats and optimization opportunities. The chapter, overall, focuses on strict, quantitative imageries of scaling regimes and is the trade-offs of measurement design, classical processing, and architectural constraints with the recent trends of the research.

2. Introduction

The issue of resource scaling is an inevitable part when trying to reconcile an idealized view of quantum advantages to concrete quantum algorithms; however, it is hard to fathom in which place it is more true than quantum kernel methods of supervised learning. Quantum kernels are inner product of data-dependent quantum feature states, which hold the promise of the high-dimensional features of the Hilbert space features that can not be obtained in the classical world. However, the promise is facing the reality

of finite sampling, noisy gates and lack of coherence. Computing the entries of kernel matrices involves preparation and measurement of quantum circuits in repeated procedures to compute the distributions of overlaps between pairs of results of the circuit. These estimations are statistical and have a sensitive dependence on the number of measurement shots and the variance of estimators; the latter is regulated by design of the circuits, selection that can be observed, and the basis of measurement. The laws of scaling between numbers of qubits, the complexity of the encoding circuit and the number of samples to a target error of estimation can thus be central to determining whether quantum kernel methods can be beneficial to practice. This chapter builds a description of these scaling relations in some detail making use of theoretical bound and modern measurement schemes that reformulates the estimation problem based on the resource's worldview.

3. Quantum kernel methods: formalism and measurement primitives

Quantum kernel machine Quantum kernel machines represent classical inputs as quantum states through a data-encoding unitary or feature map and train some overlap, or fidelity-like number between pairs of states is used to compute their kernel values. Particularly, a feature map is given map $\phi(x)$ implemented by a unitary $U(x)$ acting on an initial state (typically $|0\rangle^{\otimes n}$), the quantum kernel for inputs x and x' is $k(x, x') = |\langle U(x)U(x')^\dagger |0\rangle|^2$ or more generally $\langle \phi(x) | \phi(x') \rangle$ where $\phi(x)$ is the state prepared by the feature map. Evaluating $k(x, x')$ on hardware will involve the estimation of values of expectation or probability-tasks, which are normally achieved through repeated qualification of the quantum circuit and then measurements in relevant bases. Measuring state overlaps Directly both Starting with the simplest measurement primitive, the swap test or controlled-swap variants, cost measurement circuit depth and controlled operation whose cost can be prohibitive on near-term devices. Other methods are based on the preparation of composite circuits which include the adjoint and then measuring in computational or Pauli bases; as one example, the fidelity can be estimated by projective measuring onto $|0\rangle$ the unitary of the combined circuit $(U(x)U(x')^\dagger)$ leads to the reduction of the estimation of the overlaps to one-outcome estimation of probability. All these primitives have different resource demands: the swap test has ancilla demand and controlled-swap gate demand that are depth-dependent on the size of the system and adjoint circuits encode the two encodings sequentially and transfer resource cost to circuit depth and gate fidelity. Imperatively the statistical variance of the estimator due to the kernel is determined by whether the measurement is observing a binary event (projector onto $|0\rangle$ (required to measure)) or some higher-dimensional measuring

observable that can be written as a sum of Pauli terms; Pauli decomposition Propagates this dependence onto the nature of the feature map design and encoding of the problem.

3. Shot noise, estimator variance, and sample complexity

Any practical assertion concerning the degree of resource scaling needs to resolve the shot noise: any repeated implementation of quantum circuit under normal sampling generates a justification of an estimator with its dispersion falling as the quantity of shots. In the case of a binary outcome measurement where the probability of success is p , the variance of an empirical estimator of \hat{p} is $p(1-p)/N$ where N is the number of repeated shots. When based on kernel entries right to such probabilities, obtainment of the wanted additive error ϵ in the kernel estimate, causes a canonical scaling ($N=O(p(1-p),^{-2})$). Nonetheless, quantum machine learning estimation of a kernel is often not that straightforward: kernels can be calculated by estimators that involve a combination of a number of measured quantities, each associated with its own variance, and classical post processing can often further increase estimation error. This, therefore, needs more sophisticated sample complexity studies. These analyses take into account the worst-case or averages case variance given the input distributions, concentration properties of the process of the kernel matrices given data ensembles as well as whether estimators are unbiased and have finite or bounded variance. However, it is also important to note that there are classes of random and hardware-motivated feature maps with a potentially exponentially growing variance, and hence an exponentially growing number of shots is also required to keep additive error constant, unscalable to large scales. It is the knowledge of when this exponential variance takes place and when regimes are in which the variance is harmless (e.g., polynomials bounded with qubit number) that is of central importance in judging the feasibility of the resources required of quantum learning based on a kernel.

4. Concentration phenomena and implications for scaling

More theory Recent theory recorded concentration phenomenon in which the values of quantum kernels are concentrated around typical values in the limit of increasing qubits (or in the limit of feature maps being expressible random circuits) [9,33-35]. This non-linearity of concentration means that roughly different inputs can end up with almost the same kernel value with a large probability, when the measurement is performed with a resource of polynomial type, merging the expressivity of the used kernel and resulting in trivial classifiers. Resourcewise, the concentration effects require a tradeoff of either constructing feature maps that would not be concentrated by maintaining an underlying substantial input-sensitive structure, or spending exponentially many samples (or

exponentially accurate measurements) to cope with small variations around the concentration point. The discovery of measure of expressibility of the feature maps and the expression of requirements in the circumstances under which the concentration occurs thus have been turned into research goals. Design advice Surgeries Structured, feature maps that are problem-conscious, aware of the locality and symmetries that maintain variance in specified observables, and hybrid classical-quantum preprocessing to reduce embedding dimensions are many of the features of effective design design strategies are structural. The theoretical interpretation of concentration is frequently given by way of random matrix theory and high-dimensional probability: at exponentially growing dimensions of the Hilbert space, inner products of random high-dimensional vectors are concentrated about zero, in normalized space, and perturbations in small circuit components (or by noise causing depolarization) can only exaggerate it. Scaling thresholds designed to ensure that the estimation of the kernel at high scalings require significant investments of measurement can be determined as a result of quantitative results relating circuit depth, qubit count and the probability distribution of the values of the kernel itself.

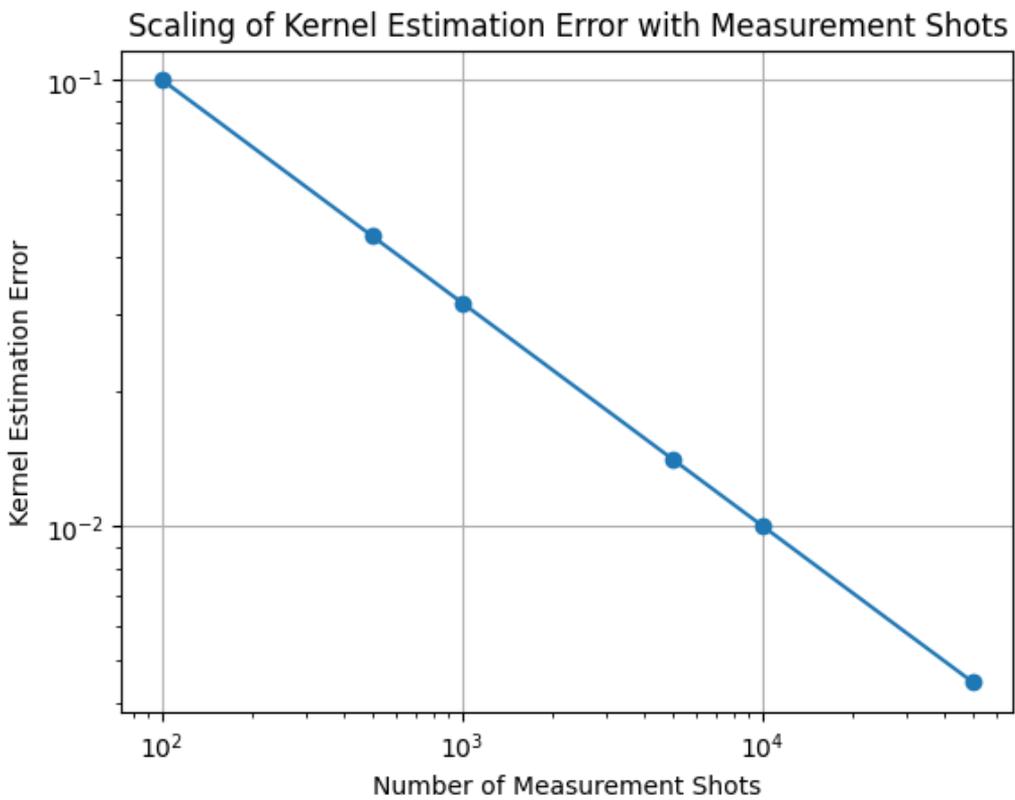


Fig 1: Scaling of Kernel Estimation Error with Number of Measurements (Shots)

5. Measurement protocols beyond single-basis sampling: classical shadows and randomized measurements

Measurement problem The classical shadow formalism and its latter derivatives, have reconstituted the measurement problem demonstrating the existence of one randomized measurement strategy, with efficient classical postprocessing, capable of providing accurate estimates to a large family of observables. The classical-shadow method puts naive repeated measurement of certain observables replacing it with randomized unitaries and measurement in a fixed basis; the set of randomized approaches constructs a compact classical model (the shadow) that can be used to estimate numerous observables simultaneously. The complexity of the sample required to estimate M observables to additive error ϵ polylogarithmically in M in the many practical case of observables families; and is based on operator norms or variance proxies. The method provides a significant saving of resources in a situation where a large number of kernel entries or a large number of expectation values require the use of the same underlying state ensemble. In addition, locally-entangled measurements extensions and shallow-circuit randomized protocols have shown that measurement ensemble modification can substitute circuit depth by a lower sample complexity with operators of some classes. This flexibility has implications of quantum kernel estimation since the amount of entries in a kernel matrix can be quadratic in number, and a scheme that is able to amortize the cost of measurement among entries is of great interest. Classical shadows do have certain drawbacks, however; First, they need the observables of interest to be limited in their variance of the chosen randomized ensemble; and they can entail complicated classical memory or computation with naive application to large datasets. Moreover, it can be seen that in particular regimes adaptive or joint measurement schemes may vastly outperform independent randomized measurements, so one should not dismiss the importance of careful consideration of the trade-offs between the complexity of measurement circuitry and sample efficiency.

6. Entangled and joint measurements: trade-offs and advances

Although independent single-copy randomized measurements are experimentally convenient, joint or entangled measurements of multiple copies of a state on the sample and entangled-base measurements can make the sample complexity reduce in selected activities. An unusual collection of protocols known as entangled measurements (define a scheme of Bell-basis two-qubit measurements) through entangled measurements, or, more generally, multipartite GHZ-based schemes of randomized measurements, can yield a better scaling of sample complexity with operator weight in most instances. The main motivation is that measurement bases that entangle can study properties of the state observable worldly more directly than measurements made of local products and that,

therefore, the sample sizes required to estimate some many-body observables are reduced. But more complex hardware capabilities are required in entangled measurements and include the ability to entangle measurement ancillas with the system, to rotate multiple qubits with non-local qubits and coheres around larger subsystems during measurement preparation. Recent research has described families of measurement parameterized by the size of entanglements, and their regimes in which entangled measurements are strictly beneficial; classes of observables in which entanglement is of no advantage and in fact imposes incompatibilities have also been explained. By experiment, it has been proposed and demonstrated (on small scale hardware) that small entanglement (e.g., Bell measurements between qubits adjacent to each other) can be used in the near future, and may can provide a practical kernel estimation improvement when pipelines are designed to take advantage of it.

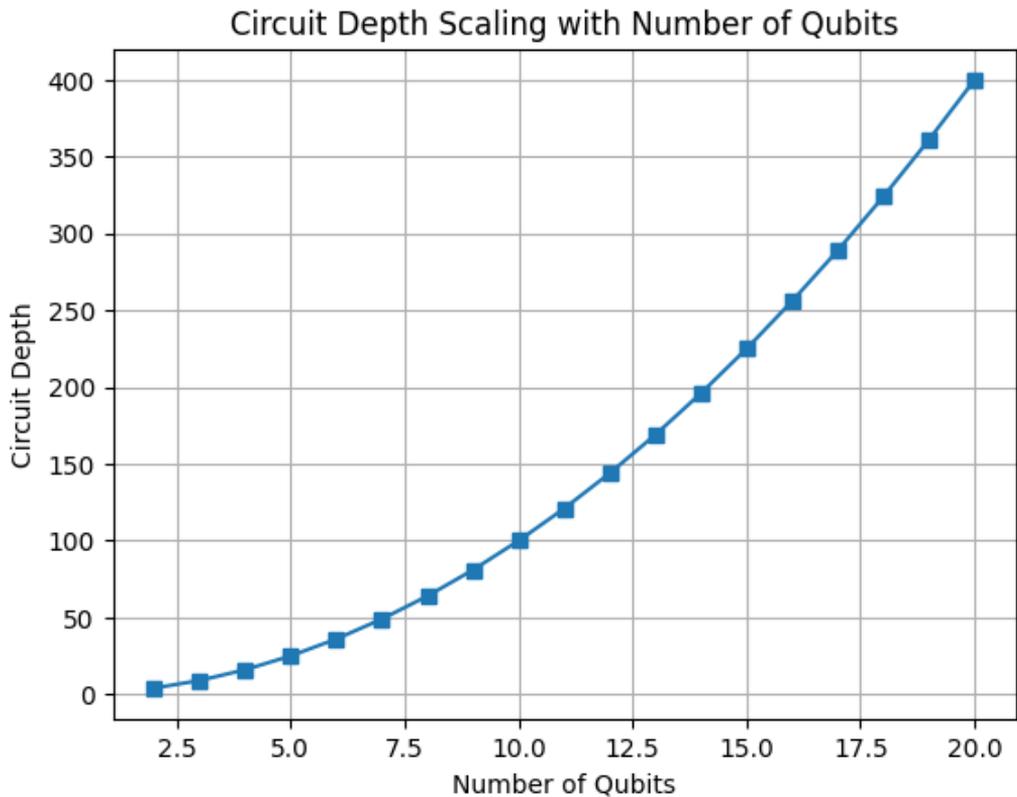


Fig 2: Resource Growth with Number of Qubits and Circuit Depth

7. Adaptive measurement strategies and estimators

Measurement strategies that are adaptive (where the decision on what to measure on later rounds), provide another dimension of improvement in the efficiency of the sample. Adaptive protocols are able to focus the measuring resources on the informative part of the observed space and minimize worst-case variance and robust to the makeup of adversarial inputs [36-38]. Adaptive strategies however make an analysis complex since they do not allow independent analysis of the measurement samples on which various proofs of concentration depend. Certain essential constraints are present as well: it is recently shown that in adversarial contexts adaptive protocols can not necessarily achieve the sample efficiencies of non-adaptive randomized protocols when applied to large classes of observables, whereas in other restricted classes of observables adaptive protocols do achieve provable advantages. In quantum kernel estimation, adaptivity may be adaptive selection of custom measurement dissimilarity of pairs or groups of inputs which are hard to differentiate, or provision of more shots to kernel entries with the highest loss to a downstream classifier. To come up with unbiased but low-variance adaptive estimators, one must pay close attention to the statistical interdependences that adaptivity brings and the guarantee that the classical postprocessing can be done on datasets of realistic size.

8. Error, noise, and mitigation overheads

No resource scaling analysis can be complete without the effects of noise being considered. Estimator variance, bias measurements, and in certain instances cause systematic drift in kernel values are all increased by gate errors, readout errors, decoherence and imperfection in state preparation. Veridical accounting of resource scaling will then have to involve overheating of error mitigation and in the case of fault-tolerant regimes, error-correction. Reducing bias can be achieved at the price of increased variance and extra classical processing by error mitigation techniques, including zero-noise extrapolation, probabilistic error cancellation, symmetry-based-postselection, and error inversion with measurements. The overall effect of mitigation is usually that the effective number of shots needed to hit an error has increased by constant or even exponential factors; in the extreme case mitigation can increase exponentially the demands in resources imposed by hitting a target error. A different resource calculus is the fault-tolerant kernels resource calculus, which additionally models error-reduced generation using physical qubits and gates that increase polynomials: or worse due to logical error rate targets - but without sampling bias due to coherent errors. The effect of the trade-off of mitigation and correction of errors in the near term and the correction of errors in the long term leads to a bimodal frontier of resources in which various areas of the scaling curves apply, respectively, with respect to maturity of hardware.

9. Classical postprocessing and computational scaling

The scaling of resources does not entirely depend on the number of quantum shots and the qubits: the role of classical postprocessing is predominant and may act as a bottleneck in the shadow-like strategies. The classical shadows need to build inverse channels and matrix multiply / spectral decomposition in reconstructing some observables and, due to the large number of kernel entries in larger datasets, building and using kernel matrices (e.g. to train a support-vector machine or to solve a large linear system to compute the kernel ridge regression or to other similar tasks) may be a substantial cost factor. Pragmatic pipeline Pragmatic pipelines are thus an attempt to trade off quantum measurements amortization (measure once, estimate many quantities) and practical approximations based on classical features-low-rank kernel approximations, Nystrom methods, random feature analogues, and sparse inference tools, in order to make the overall system manageable. Also, classical shadowing and online versions of the updates to the kernel have been suggested in order to restrict the growth of memory. The selection of classical algorithm and numerical precision requirements is further passed on to the accuracy desired of quantum kernel estimates and hence the investment of quantum shots per kernel entry.

10. Practical guidelines for resource allocation

Putting the theoretical knowledge and practical limitations together allows us to obtain a number of practical-minded rules of allocating resources in quantum kernel estimation processes. To start with, designs of feature maps should first promote input-dependent variance in low-weight observables, and where it allows locality or structured encodings should be preferred when using near-term hardware. Second, randomized measurement ensembles can be amortized across several kernel entries, and the classical shadows should be used or shallow randomized circuits should be used where possible. Third, in places where hardware allows, one can improve scaling in observable families thanks to limited entangled measurement bases (i.e. two-qubit Bell measurements among interacting qubits). Fourth, explicitly compute error mitigation overheads during translation of target estimator precision to shot budgets; compute budget at least pointing at poly growth of such overheads of typical mitigation strategies on NISQ devices. Lastly, reserve classical approximations, e.g. low-rank factorizations, Nystrom approximations, sparse solvers, etc. early on in the pipeline to ensure that quantum measurement accuracy counts up with the classical model's capacity to utilize the available accuracy. Such guidelines focus on co-design in feature maps, circuits, measurement plans and classical postprocessing to turn out resource-efficient workflows.

11. Two comprehensive summary tables

Table A: Measurement Protocols — Resource characteristics and scaling

Protocol	Measurement type	Typical resource scaling (shots, qubits, depth)	Strengths and limitations	Practical implementation notes
Projective computational-basis measurement of adjoint circuit	Single-copy projector onto	$ 0\rangle$ after $U^\dagger(x)U(x')$	Shots: $O(\epsilon^{-2})$, Qubits: n , Depth: $\text{depth}(U(x)) + \text{depth}(U(x'))$	Simple to implement; depth doubles; variance depends on state overlap
Swap test / controlled-swap	Joint two-copy measurement with ancilla	Shots: $O(\epsilon^{-2})$, Qubits: $2n+1$, Depth: ancilla-controlled-swaps	Direct overlap estimator; ancilla and controlled operations costly	Requires high-fidelity controlled-SWAP; overhead in qubit connectivity
Pauli-decomposition estimators	Decompose projector into Pauli terms measured via randomized Pauli bases	Shots: $O(\sum \text{variances} / \epsilon^2)$, Qubits: n , Depth: shallow per randomized basis	Flexible; connects to classical shadows; variance depends on Pauli weight	Good when observable has low Pauli weight or locality
Classical shadows (single-qubit Clifford)	Randomize product unitaries, measure computational basis	Shots: $O(\log(M) B / \epsilon^2)$ for M observables with variance proxy B	Amortizes across many observables; efficient classical postprocessing	Postprocessing cost grows with M ; best for many-query regimes

Locally-entangled shadows (Bell, small blocks)	Randomized measurements with entangled block unitaries	Shots: Improved scaling for certain operator classes; Qubits: n , Depth: moderate	Can reduce Pauli-weight dependence; hardware-costly vs. product shadows	Tradeoff entanglement size against hardware feasibility
Global Clifford / deep randomized ensembles	Global randomization approximating Haar	Shots: Optimal for low-rank observables; Depth: large	Best sample complexity for global properties; deep circuits required	Feasible only on error-corrected or highly coherent devices
Adaptive measurement schemes	Measurements chosen based on prior outcomes	Shots: can reduce worst-case sample complexity	Efficient for adversarial or focused tasks; complex analysis	Requires classical control loop and careful bias correction

Table B: Resource trade-offs for kernel estimation primitives

Primitive	Quantum cost (per pair)	Classical cost	Sensitivity to noise	Best-suited regimes	Scalability concerns
Direct fidelity via adjoint circuit	Depth: $O(\text{depth}(U(x)) + \text{depth}(U(x')))$, Shots: $O(\epsilon^{-2})$	Minimal per-pair; kernel matrix assembly $O(N^2)$	Moderate; coherent errors bias estimates	Small to medium datasets; low-depth feature maps	Quadratic number of pairs leads to large shot budgets
Swap-test overlap	Qubits: $2n+1$; controlled ops	Minimal; but ancilla handling	High sensitivity to controlled-gate error	Small n , when direct adjoint not available	Doubles system qubits; connectivity

					constraints
Shadow-based many-pair estimation	Single-shot per circuit reused across pairs then postprocessed	Postprocessing $O(M)$ or matrix operations	Mitigation needed for readout and unitary errors	Large datasets where kernels require many entries	Classical memory and compute for shadows can become bottleneck
Compressed / low-rank quantum kernel approximation	Fewer pairs estimated, interpolation	Classical low-rank solvers; Nyström	Less sensitive if low-rank valid	Large datasets with low intrinsic dimension	Approximation error must be controlled
Entangled - measurement aided estimation	Similar shot reductions for some observables	Moderate; complexity in estimator	Entanglement amplifies decoherence sensitivity	Intermediate-scale devices with moderate entangling capability	Hardware-specific; not yet broadly available
Fault-tolerant logical kernel estimation	Very large physical qubit overhead, low logical error	Classical costs standard; solvers scale with data	Logical-level robustness to noise	Long-term, fault-tolerant advantage scenarios	Requires huge physical qubit counts; long horizon

12. Emerging approaches and research frontiers

Various research frontiers have potentials to alter the existing concept of resource scaling to resource estimation of the kernels. The systematic exploitation of (optimized) entangled randomized measurements, designed to be optimized with the observed quantities of interest, is one such direction; and theory indicates that entangled blocks

noted in several modest sizes can scale a large fraction of Pauli-weighty observables, and be implemented in near-hardware hardware technologies. The other developing field is the design of estimators, in which the incomplete U-statistics and cyclic permutation technique can be used to refine the design of estimators to reduce the effective variance of the kernel estimators and utilize any structure present in the data. Concurrently schema Adaptive hybrid classical-quantum allotment plans that either incorporate the learning goal (downstream) (such as kernel-target fit or classifier margin) in shot distribution choices are also actively being developed. On the hardware front, more capability of measuring and resetting in the middle of the circuit, and more realistic two-qubit gates, as well as better qubit connectivity, will further increase the range of measurement ensembles that are practical. Lastly, the benefits of theoretical progress that understand the environments in which the values of the kernel become concentrated or not will serve to determine the demarcation of the boundary of the classically simulable and quantum-advantaged regimes.

13. Limitations, open problems, and cautionary perspectives

Various research frontiers have potentials to alter the existing concept of resource scaling to resource estimation of the kernels. The systematic exploitation of (optimized) entangled randomized measurements, designed to be optimized with the observed quantities of interest, is one such direction; and theory indicates that entangled blocks noted in several modest sizes can scale a large fraction of Pauli-weighty observables, and be implemented in near-hardware hardware technologies. The other developing field is the design of estimators, in which the incomplete U-statistics and cyclic permutation technique can be used to refine the design of estimators to reduce the effective variance of the kernel estimators and utilize any structure present in the data. Concurrently schema Adaptive hybrid classical-quantum allotment plans that either incorporate the learning goal (downstream) (such as kernel-target fit or classifier margin) in shot distribution choices are also actively being developed. On the hardware front, more capability of measuring and resetting in the middle of the circuit, and more realistic two-qubit gates, as well as better qubit connectivity, will further increase the range of measurement ensembles that are practical. Lastly, the benefits of theoretical progress that understand the environments in which the values of the kernel become concentrated or not will serve to determine the demarcation of the boundary of the classically simulable and quantum-advantaged regimes.

14. Conclusion

Scaling resources in quantum kernel estimation and measurement protocols is a complex issue and should be optimized concurrently in quantum circuit design and measurement strategies, performance of countermeasures to error, and in classical processing. Quantum kernel methods are empirically viable based on the well-crafted engineering decisions that maintain the input-dependent variance, exploit the measurement amortization with randomized impressed ensembles or shadows, and make considered use of the available classic computations. Although recent developments in classical shadow tomography, entangled randomized measurements, and adaptive allocation algorithms do open up possibilities to ease up the load on measurement, such underlying limitations as concentration effects and noise effects creating variances still are highly influential factors that dictate eventual scalability. Theoretical conditions that restrict the complexity of the samples of realistic noise models, and experimental proof that resource trade-offs follow on existing hardware will be necessary to help to the boundaries of what the future of kernel-based quantum machine learning will be.

Chapter 6: Sample Complexity and Generalization Bounds in Quantum Learning

1 Abstract

This chapter also summarizes the modern theoretical context of complexities of samples and generalization limitations in quantum learning theory and incorporates classical-learning-theoretic principles with the operational limitations peculiar to quantum data, quantum hypothesis classes and quantum-measuring apparatuses of the learning process. It is aimed at providing a coherent description of the definition, measure and bounds of sample complexity in quantum regimes along with an explanation of the main mathematical methods including analogies to VC-dimension, Rademacher complexities, covering numbers, information-theoretic and hypothesis-class-specific lower bounds as well as to describe state of art expectations and current problems that have been encountered over the past half-decade of active research. I highlight conceptual distinctions that are created due to measurement back-action, noncommutativity, entanglement and the hybrid classical-quantum setups characteristic of systems in the near term, and relate these distinctions to actual sample complexity behaviors of canonical problems including, but not limited to, learning quantum states, quantum channels, quantum measurements and quantum-kernel-based classification. The rigorous findings that exemplify representative limitations and opportunities in quantum sample efficiency are defined and put into context in relation to the modern literature and I conclude the work by predicting future research directions which seem especially promising in minimizing effective sample demands or in estimating generalization behaviour resilient to hardware noise and finite scale effects.

2. Introduction

The issue of the number of data points needed to learn a predictive model is the main theme in the statistical learning theory, and when the data, models or the learning processes involve quantum mechanics it becomes a question of new dimensions since

quantum information is brittle, moreover it is utterly unlike classical information. Sample complexity results, upper and lower bounds relating the number of training examples to target accuracy and target confidence, in classical learning theory are regularly stated in terms of combinatoric quantities or VC-dimension, Rademacher complexity, covering numbers or Metric entropy of hypothesis classes and they depend on copying inequalities that rely on the assumption of independent sampling. There are three different resources to be carefully differentiated in the quantum setting, namely (i) copies of quantum states (or quantum examples) following an unknown distribution, (ii) quantum queries on unknown channels or oracles, and (iii) classical labeled examples that are the result of measurements on quantum states or quantum-processed data. The fact that there is operational difference between many independent copies of quantum states and having access to arbitrary queries to a quantum oracle is the cause of the difference of measures of complexity and subtle separations/equivalences between quantum and classical sample complexities in particular models. The chapter decomposed these regimes, clarified the main types of theorems that have been established (upper bounds using constructive algorithms and lower bounds using information-theoretic or adversarial constructions) and highlights where quantum effects like the effects of measurement-induced collapse and noncommutativity have fundamentally changed the analysis and conclusions. Some recent survey and technical papers have built upon previous foundations with the demonstration not just that in most canonical settings quantum and classical sample complexities are equalized at scales of constants but also that there are indeed practically significant quantum learning tasks in which quantum sample complexity can be proven to be larger or which demand qualitatively new algorithmic patterns.

3. Preliminaries and formal models

There is need to commence with accurate definitions of the learning issues to be compared. The Probably Approximately Correct (PAC) learning model applied to quantum objects is one model that is studied widely: a family of quantum states or channels or measurement operators, and a learner is given either classical labels of these states (generated by the process of measurement), or a quantum copy of actually quantum states sampled according to some unknown distribution [3,39-41]. In the case where the learner is provided with quantum examples, the resource unit is a copy of a quantum state selected according to the distribution that is actually the indivisible copy, and any measurement made can permanently modify it. In the case of channel learning, it is possible to think about either access to multiple roles of the unknown channel (sample complexity in terms of number of queries) or access to pairs between inputs and outputs of the channel operating on distributed inputs. Another and equally significant model is that of agnostic learning where no presupposition is made that the data

completely concurs with any of the hypotheses in the category. The aim of the models is often to discover a hypothesis which has expected loss over the underlying distribution within ϵ of the optimal hypothesis in the class, and whose probability $1-\delta$; sample complexity is the minimum number of examples (or queries) needed to ensure this. The quantum setup demands extra treatment since the measurements of the learner can be adaptive, and making joint measurements on many copies can be beneficial or exhibit inherent drawbacks being due to noncommutativity. Formalizations in the literature have made it clear that most classical definitions have been carried over with the addition of operational restrictions, as well as information-theoretic measures of the distinguishability of quantum objects, trace distance, fidelity and distances of the same type, which show up explicitly in both bounds and impossibility theorems.

4. Key technical tools: quantum analogues of classical complexity measures

Classical classicality metric translation to the quantum world has been conceptually easy in certain managed and notionally challenging situations. In the case of classes of quantum measurements or quantum states, it is possible to define VC-like combinations of classifiable labelings that can be generated by a fixed set of inputs by a measurement subject to a restricted set of measurements; results in known sample complexity bounds have constant or logarithmic corrections depending on quantum measuring structure. The coverage by balls based on the trace norm or diamond norm, as were developed in metric entropy and covering-number methods, are still exceedingly helpful: when one can cover the hypothesis space (i.e. a set of density matrices, POVMs or unitary channels) with balls, then it is possible to bound uniform deviations between empirical estimates and the true expectations, and in the process regulate generalization error. A quantum analogue of Rademacher complexity can be described by using random sign functions in place of randomized measurement interactions or the hypothesis space being modeled by parameterized quantum circuits and fluctuations of the empirical loss as an empirical loss are analyzed in a randomized labelling model. Inequalities of concentration on quantum observables, noncommutative analogues of classical Bernstein or Hoeffding inequalities, are irreplaceable since some of the most common bounds are in terms of tail estimates on sums of independent measurement outcomes; they are usually subject to operator norm and commutation hypotheses and are currently actively developed mathematically. Besides these classical theory extensions, purely quantum tools can be found in sample-complexity. quantum tomography lower bounds (a count of copies of states it takes to learn an unknown density matrix to statistical accuracy) give minimal impossibility bounds, quantum information measures (such as Holevo information or accessible information) are used to derive lower bounds on the number of samples to distinguish amongst a set of candidate states or to learn distributional properties of quantum ensembles. A combination of these tools of analysis

allows researchers to discover upper bounds with constructive protocols, which usually include shadow tomography and randomized measurements or collective measurements, and discover lower bounds by information-theoretic arguments that quantify the shared information of samples and unknown parameters.

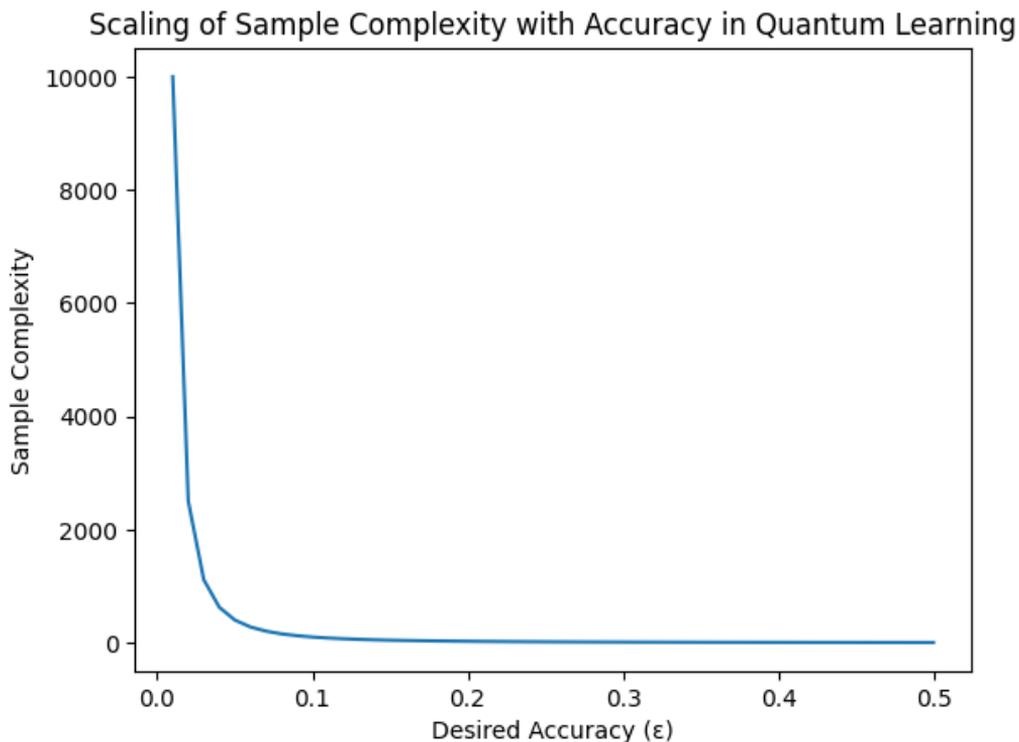


Fig 1: Sample Complexity vs Desired Accuracy (ϵ)

5. Classical versus quantum sample complexity: equivalences and separations

One of the focal ideas of the quantum learning theory is when quantum resources have a strict reduction of sample complexity, when quantum resources have a constant factor improvement, and when quantum resources do not have any asymptotic gains. Historical findings showed that in some instances of PAC and agnostic learning environments the sample complexity of quantum learners is identical to the sample complexity of classical learners, up to constant factors and this implies that having access to quantum examples or quantum queries does not necessarily provide quantum learners with asymptotic advantages in the required number of examples in order to accomplish a specified generalization accuracy. Less dramatic studies on the other have shown regimes in which quantum sample complexity is either significantly worse or in which various resource trade-offs become apparent: tasks where noncommutativity of measurements requires a learner to use many copies, or where noncommuting measurements by over many copies

can be better than separable measurements can. In the more recent research, these separations have been refined by studying particular types of classes- stabilizer states, or low-rank states, or families of channels parameterized directly by the noise- or NISQ-era imperfections and also quantifying the impact of these imperfections on sample-complexity trade-offs. The resultant implication is that even though asymptotic quantum advantages in raw sample complexity are uncommon to broad concept classes, domain-structured tasks, model-conditioned prior information, and intelligent measurement schemes, sample-efficiency advantages are possible in practice in domain-structured tasks. This subtle image, typically backed by constructive algorithms and closely matching lower bounds in numerous situations, implies that pursuit of the practically useful quantum learning benefits is most appropriately done through identification of structural priors or through creation of the hybrid learning structures incorporating both the quantum state structure and the classical inductive biases.

6. Representative upper bounds: constructive algorithms and their sample counts

Quantum learning literature is likely to provide concrete upper bounds related through constructive algorithms [36,42-44]. One such prominent tool is based on classical shadows tomography and random lackadaisical measurement plans that encode information on an unknown quantum state in a tiny set of classical shadows whose purpose is to estimate myriads of expectation values with provable accuracy; in case the learning target is constrained to a family of observables with fixed controls, shadow-based methods can achieve desirable sample complexities growing polylogarithmic in the number of observables of interest, but growing polynomially in inverse desired accuracy. To achieve a sample complexity estimate, learning parametrized quantum circuits, or kernel models, constraints on practice of parameterized families (in terms of covering numbers or parameter norm constraints) are commonly derived relative to the differences between the empirical risk estimators associated with successive measurements. Operating in the PAC-style learning of quantum measurement classes, noncommutivity-aware, sample complexity bounding, approaches that optimize over sampler-constrained partitions of POVM elements have been proposed; it is shown how sample-efficient algorithms may be obtained in some circumstances by partitioning the elements of a POVM into jointly-measurable subsets. Specialized algorithms, which use structural priors: sparsity, locality, or low entanglement, can be designed so that samples necessary to generate specialized quantum states or states at a given phase of matter scale with a small scale factor; e.g. polynomially with rank or locality instead of the full dimension of the Hilbert space. These are constructive algorithmic upper limits, which can be used to inform much realistic protocol design but their assumptions, such as the availability of collective measurements, noiseless gates or structured priors need to be carefully calibrated against the facts of experiments.

Table 1 — Sample Complexity Results for Representative Quantum Learning Problems

Problem / Task	Access Model	Typical Sample Complexity (scaling)	Key Assumptions
Full arbitrary n-qubit state tomography	Independent copies of unknown state	Exponential in n ($\Theta(2^n)$ up to polynomial factors)	No structure assumed; global fidelity/trace-norm accuracy
Low-rank (rank-r) state learning	Independent copies; possibly collective measurements	Polynomial in n and r (e.g., $\text{poly}(n, r, 1/\epsilon)$)	State is low-rank; access to suitable randomized measurements
Learning quantum channels (general CPTP)	Channel queries / uses	Scales with dimension^2 ; often exponential in n for worst-case	No structure; diamond-norm accuracy
Learning unitary channels (parametrized families)	Channel queries	Polynomial in parameter count and $1/\epsilon$ for structured families	Assumes low-dimensional parametrization or sparsity
PAC-learning finite quantum concept class C	Quantum or classical examples	$O(\log C)$	C
Learning stabilizer states with noise	Quantum examples with noise	Can be hard; sample complexity grows with noise parameters; often polynomial when noiseless	Structure: stabilizer formalism; presence/absence of noise critical
Measurement (POVM) learning	Quantum examples + measurement outcomes	Bounds depend on operator-class combinatorics; often $O(\text{dimension-dependent} / \epsilon^2)$	Partitioning jointly-measurable subsets; POVM class structure
Quantum kernel methods (training classical predictor)	Classical labels from quantum feature evaluations	Sample complexity tied to kernel effective rank / eigenvalue decay; can scale as $O(\text{Tr}(K)/\epsilon^2)$	Kernel eigen-decay assumptions; classical generalization theory applies
Classical shadows for many observable estimation	Randomized single-copy measurements	$O(\log M / \epsilon^2)$ to estimate M observables simultaneously	Observables bounded; randomized measurement ensemble
Quantum Boltzmann machine parameter learning	Samples of expectation values / classical-quantum hybrid data	Polynomial in system size under structural assumptions; otherwise challenging	Relative-entropy losses and model structure mitigate barren plateaus
Learning ground-state properties in gapped phases	Samples from states within phase	Polynomial in n for certain property learning under phase assumptions	Assumes phase locality and access to sample states from same phase

Shadow-based tomography for expectation estimation under noise	Randomized measurements with error mitigation	Sample complexity increases depending on noise but retains log M dependence under mild noise models	Noise model assumptions; mitigations applied to empirical estimates
--	---	---	---

7. Bases of information theory and adversary dialogue Representation lower bounds.

The standard formulation of lower bounds involves the construction of ensembles of quantum objects (states or channels) that are hard to distinguish in pairs (and using a restricted set of protocols) and by demonstrating that an algorithm that successfully solves a given problem can extract a quantity of information to the samples (equally random), one can then provide a lower bound on sample sizes by establishing a bound on the mutual information or accessible information between samples and unknown labels. As approaches to quantum measurement Classical models like the inequality of Fano, the method of Le Cam, have quantum versions which make use of quantum information measures, and give strong strict or near-strong lower bounds to problems, including full-state tomography, where sample complexity grows exponentially with the size of the system in the absence of structural priors. Changes in More recent lifting methods express the query complexity lower bounds of an access model in terms of the sample complexity lower bounds of that same access model, giving results indicating how lower bounds of one access model translate to those of another. Other recent contributions of importance have also studied the complexity of samples with noise and show how noisy measurements or noisy preparations of state can exponentially amplify the numbers of samples required, or even make certain regimes of parameters unlearnable at all. To the extent that results based on lower bound are sufficient to warn about over-optimistic concepts of sample efficiency of all universal quantized problems, and to constrain algorithm designers to one choice of restricting algorithm classes or to capitalize on problem-specific structure in realistically learned situations.

8. Generalization bounds for parametrized quantum models (QNNs and quantum kernels)

Generalization The performance of a learned model on unseen data (also known as generalization) has recently received focused attention to parameterized quantum circuits (also known as quantum neural networks, QNNs) as well as to quantum kernel methods. In the case of QNNs, analyses of generalization refine classical analysis of uniform convergence to arguments, finding a bound on the capacity of the circuit family,

with respect to the number of gates, the number of layers, the norm of parameters, and the effective Lipschitz constants of induced measure expectation depending on parameters. Recent literature has demonstrated that deep, unstructured quantum circuits are subject to two conflicting forces on the one hand, deep circuits are capable of expressing complex functions and hence have high representational capacity; and on the other hand, expressivity can lead to barren plateaus in training landscapes, and expressivity can result in empirical risk estimators with high variance, which can both make deep circuits more vulnerable to scale complexity, and reduce the generalization strength of deep circuits in the presence of finite-sample noise. The quantum kernel methods offer a different approach in which a quantum feature map calculated by a quantum circuit and the learning takes place classically on kernel matrices; in this case, sample complexities can be characterized by uniform stability bounds and generalization bounds based on kernel norm and structure of the kernel (effective rank, eigenvalue decay, and concentration) determines the structure of the quantum feature map. According to recent theoretical and empirical studies, whereas certain quantum kernels may enhance the feature separability, the overall feature separability benefits may disappear with overly strong concentration of the kernel matrix, or with increasing effective dimensionality with qubit count and, as a result, demand increasingly more data to prevent overfitting. Prescriptive advice on the design and training of kernels have started being given by systematic studies relating eigenvalue decay to finite-sample generalization by quantum kernels to kernel alignment.

9. The role of noise and NISQ-era constraints

The current wave of research poses a dramatic inquiry involving how the elements of noise and device constraints, the characterizing traits of the NISQ era, impact the complexity of samples and the generalization in general. Noise may change the content of information of each sample and the variance of risk estimators, gate errors and readout noise may decrease the mutual information of training samples and the target parameters, leading to increased sample requirements to achieve a desired confidence level. However, on the other hand, there are implicit regularizers of noise models, that in some situations generalize better by irregularizing highly expressive hypothesis classes; it is a very subtle effect, but it can be attributed to the interactions of noise spectra with circuit architecture, as well as noise landscape. In sample complexity terms, a number of new studies are making randomized version of the degradation in distinguishability created by noise and obtain lower limits indicating that above some amount of noise, no algorithm can effectively learn a target using a sufficiently large number of samples, namely polynomials many. Parallel works investigate noise-risks protocols, such as noise-sensitive measurement bases, noise reduction used on empirical measurements, and robust editions of shadow tomography, that aim at the recovery of much of the

theoretical sample intellect of the ideal hardware. The also practical implication is that serious comparison claims about the learnability of quantum learners need to explicitly describe hardware models and bridged theory and experiment need closer interaction between modulization of device error behaviour and statistically analysing.

10. Learning states, learning channels, and learning measurements.

In order to make the above theory concrete, it is convenient to refer to the classical problem classes and the most widely known sample complexities of each of them [40,45-47]. Various Conditionally complete in general case Tomography of n qubits, without structure A sample complexity that is exponentially dependent on n is required to obtain a measurement of the full quantum state of a general layout n -qubit mixed system; nevertheless, with low-rankness or locality in the state, sample methods can be designed that are polynomially dependent on n . Quantum channel learning is generally prepared by sample/query complexity which is linear in channel dimension and accuracy of desired quantum channel mostly quantified in diamond norm; special cases such as unitary channels, sparsity-structure CPTP maps allow better sample complexity. The theory of quantum measurements POVMs Learning quantum measurements Measurement elements may be noncommitting, and such samples may be prohibitively complex; the scale problem here typically is a combinatorial partitioning into jointly-measured subsets and goes up to effective analogous scale problems in spaces of operators. In all these case studies, the theme is straightforward: structural assumptions and algorithmic ingenuity are turning haplessly generic problems into polynomials sample-complex learning problems, although the cost of eliminating those very assumptions is, in most instances, an asymptotic proscribed.

11. Active issues and identification of opportunities.

However, even in the context of quick development, there are some essential questions that are still open and are waiting to be taken into consideration. To begin with, the implication of a tighter definition of the conditions under which quantum data, in contrast to classical data on systems of quantum systems, are available is still non-existent; would it be nice to have fine-grained criteria of the sort by which structural properties of problem classes (such as symmetry, sparsity, or decoupling) entail quantum sample-efficiency gains. Second, realistic-noise-model effect on lower bounds should be expanded more treatment-definitely-operationally: besides worst-case, average-case and distributional analysis of particular experimental noise model channels would be useful in operationalising complexity-theoretic predictions. Third, no asymptotic generalization theory of QNNs should be bridged with empirical behaviour empirically

observed in experiments: there are ranges of circuit norms and number of parameters, but these quantities should be associated with eigenvalue decay in the kernel or with measures of margin in experiments, an area of study. Lastly, there exists the most promising pathway of problem reduction by algorithmic development, which distillates problem aware embeddings with structured measurements with classical regularization, showing the most promising pathway towards practical sample requirements in near term quantum learning problems. To answer these questions, quantum information will need to be cross-pollinated with statistical learning theory and numerical experimentation of realistic hardware will be required.

Table 2 — Techniques for Generalization Bounds and their Typical Scalings

Technique / Tool	Quantity Controlled	Typical Scaling Behavior	Applicability (Problem Classes)	Representative Works
VC-dimension analogues for operator classes	Uniform deviation between empirical and true risk	Sample complexity $\propto (\overline{VC} + \log(1/\delta))/\epsilon^2$	Binary classification with quantum measurement concept classes	PAC learning of measurement classes. (Proceedings of Machine Learning Research)
Covering numbers / metric entropy in trace/diamond norm	Uniform convergence over hypothesis space	Sample complexity $\propto (\log N(\epsilon))/\epsilon^2$ where $N(\epsilon)$ is cover size	State/channel estimation with norm-based error measures	Metric-entropy-based bounds in tomography. (arXiv)
Noncommutative concentration inequalities	Tail bounds for sums of observables	Exponential tails with operator-norm dependence (quantum Bernstein/Hoeffding)	Analysis of empirical risk and variance in quantum measurements	Concentration for operator sums used in generalization proofs. (Wiley Online Library)
Rademacher complexity (quantum-adapted)	Expected sup deviation of empirical processes	Sample complexity scales with empirical Rademacher complexity / ϵ^2	Parameterized quantum circuits and kernel methods	Rademacher-style analyses adapted to QNNs and kernels. (arXiv)
Information-theoretic lower bounds (Holevo/Fano)	Lower bounds on distinguishability and mutual information	Sample complexity lower bounds \propto information-deficit / \log	State discrimination, minimax lower bounds for tomography	Lower bounds via Holevo/Fano and accessible information. (arXiv)

Shadow tomography and randomized measurement theory	Estimation error for many observables	Sample complexity $O(\log M / \epsilon^2)$ for M observables	Estimating many expectation values simultaneously	Classical shadows and randomized measurement guarantees. (arXiv)
Kernel-eigenvalue / effective dimension bounds	Generalization gap for kernel ridge/regression	Sample complexity \propto effective dimension / ϵ^2 (effective dimension = $\sum \lambda_i / (\lambda_i + \lambda_{\text{reg}})$)	Quantum kernels and classical learning on kernel matrices	Quantum kernel generalization analyses. (arXiv)
Structural prior exploitation (low-rank, locality)	Sample complexity reduction via priors	Polynomial or polylog scaling in n when structure holds	Low-rank states, local Hamiltonian ground states	Learning with structural priors and phase-based algorithms. (Nature)
Noise-aware bounds and robust statistics	Bounds that include hardware noise dependence	Sample complexity increases as noise reduces SNR or increases bias	NISQ-era algorithms and noisy measurement protocols	Analyses incorporating error models and mitigation. (arXiv)
Sample-to-query lifting techniques	Relate query complexity to sample complexity	Quadratic or other lifting relations between query and sample costs	Bridging oracle-query lower bounds to sample models	Recent lifting lower bounds and their implications. (epubs.siam.org)

12. Expansive practice and experiment implications.

The theoretical image provided by the existing generalization and sample complexity analyses, has operational implications on a case-by-case basis to the experimentalist and practitioner. To start with, it highlights the importance of integrating domain knowledge into models: be it low-rank priors to learn states or higher-rank structure to classify, it is clear that structure will decrease the amount of samples needed and both remove noise and decrease sensitivity of measurements. Second, it emphasizes measurement design as a first-class engineering lever; measuring bases that give maximum information in a sample or group of observables can make sure that the constants of the sample count depend on the selection of the measurement bases; making tasks hitherto expensive to do to become practical. Third, it also highlights hybrid models, in which quantum equipment supports expressive capabilities and learners in classical fields deal with generalization, as a near-term viable course of action. Lastly, with the continued integration of noise models and finite-sample statistics into theory, a more prescriptive

recommendation of protocols all explicitly trading fidelity of devices, available copies and target accuracy should be expected. Overall, the result of the theoretical toolkit we already have today is the conversion of abstract sample complexity constraints into practical protocol decisions which can be checked with costs of hardware by experimenters.

13. Conclusion

The problem of sample complexity and generalization constraints of quantum learning has come out of isolated theoretical curiosity to a formal research agenda with practical implications to the experiments of the near term and algorithm design. The changing story is that of delicate delicacies: although universal asymptotic quantum improvements to sample complexity are uncommon, quantum-specific effects and structure-exploiting algorithms are real possibilities of quantum space to lower effective sample complexity of a wide range of learning problems of practical interest. The combination of measurement theory, noncommutative concentration, information-theoretic lower bound, and structure-aware algorithm provides a powerful tool kit in the future development direction. The ongoing independence of narrowing no asymptotic guarantees, integrating hardware sound of noise into statistical analysis, and end-to-end enhancements in the lab will be essential towards the digital clinics of theoretical hopes to dependable quantum-enhanced learning.

Chapter 7: Quantum Data Encoding: Cost and Scalability of Input Preparation

1 Introduction

In quantum algorithms, speedup over the classical computer is often assured to be asymptotic, especially in linear algebra and optimization problems, in machine learning models, as well in compressing classical data based on its formatting to a quantum system. The issue of encoding classical data into quantum states, quantum data encoding or state preparation is hence not an engineering byword but it is a key part that determines the overall cost, feasibility, and scalability of quantum solutions. The chapter studies taxonomy of encoding strategies, compares their costs in terms of resources (quantity of qubits, gates, and circuit depth). By so doing I observe both recent theoretical and experimental progress, architectural limitations (such as those due to proposed quantum random access memory, QRAM) as well as consider how useful performance can be restored with only very reduced resource cost using approximate and structure-aware encodings. The discussion will refer to both modern literature and surveys to contextualise the contemporary developments, and to both explicitly identify, which of the matters concerning claiming quantum advantage, they are conditional on optimistic assumptions regarding the overheads in preparing the inputs. In the case of claims and technical outcomes relying on the ongoing research in hardware and architectures, I refer to current surveys and papers exemplifying the conflict between algorithmic progress and system-level expenses. However, in the case of an evaluation of large-scale data loading assumptions, of particular interest is criticism of QRAM models and explanations of when active QRAM designs can degenerate asymptotic speedups.

2. The primary part in end-to-end complexity state preparation.

Any quantum algorithm which takes in classical data, should initially carry out a mapping of the bitstrings or numerical vectors into some quantumization. This may be as straightforward as putting a dataset index in to computational basis states, or even as

elaborate as amplitude encoding where N dimensional classical vector is encoded using the amplitudes of a $\log_2 N$ -qubit state. The resulting computational complexity of that mapping, in terms of the number of qubits, two-qubit and single-qubit gates, circuit depth, use of ancillary, and classical pre-processing is a direct cause of the wall-clock time of the algorithm. Theoretically split-off costs of speedup: indeed, full amplitude encoding of an arbitrary vector will typically take a number of gates which is polynomially dependent on the dimension of a vector, but typically exponentially dependent on the number of qubits with no structure in which it can be done and where the additional overhead can therefore negate in practice the quantum advantage. This has led to two opposite directions in literature: initially, to design architectures including QRAM, which tries to offer sublinear access of massively sized classical memories in superposition, second, to develop algorithmic and compiler methods to utilize structure, approximation, and compact representations (e.g., matrix product states, structure of low-degree polynomials) to synthesize useful states using much fewer resources than would be indicated by the worst-case complexity. Recent examples of encoding discretized polynomials and encoding approximations of amplitudes exemplify the second approach and prove that in the case of structured inputs state preparation schemes with polynomially decreased number of gates can be constructed, and it is possible to design shallow circuits with polynomially reduced number of gates and suitable to near-term devices.

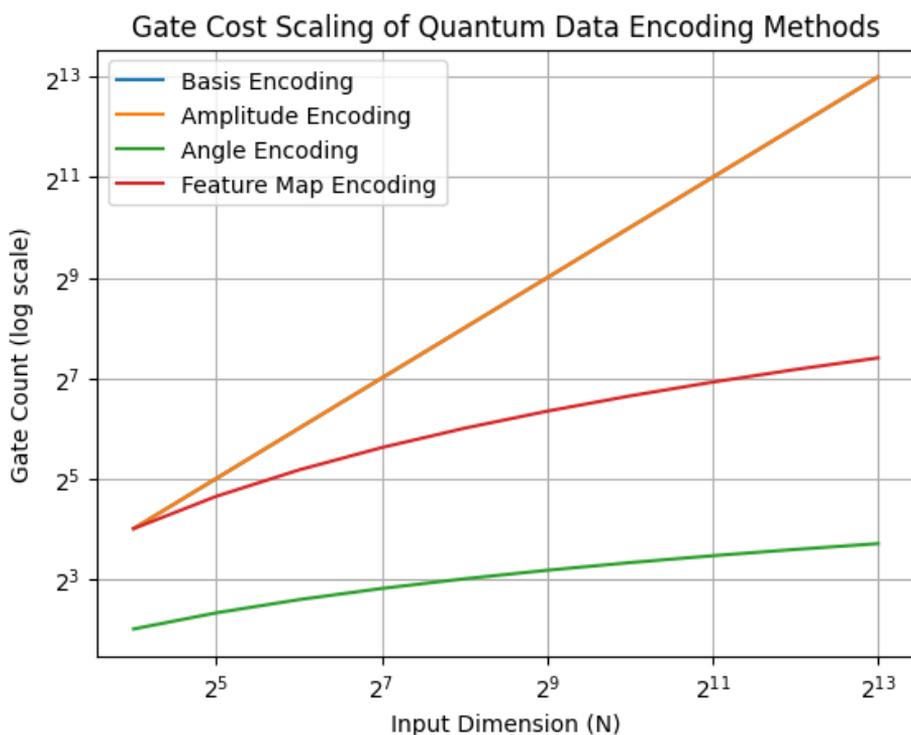


Fig 1 : Scaling of Gate Cost with Input Dimension for Different Encoding Schemes

3. Taxonomy of encoding methods and their resource scalings

Canonical encoding paradigms are several paradigms that are repeated in quantum algorithms and experiment implementations [3,48-50]. Store knowledge in the computational basis in many qubits and is therefore simple to prepare to take slim, indexed data but inefficient to store prevalent numerous vectors due to its requirement of numerous qubits to encode dimension. ~ An N dimensional classical encoding Amplitude encoding encodes the values of a N dimensional classical system, compressing it into amplitudes of $\log_2 N$ qubits and is particularly appealing due to its qubit-efficiency; however, the cost of encoding an arbitrary state is exponential in N against the size of N: that is, it cannot be reliably prepared with circuit with $\log_2 N$ qubits in an arbitrary form in a typical decomposition. Angle form or rotation-based encodings encode scalar features as rotation angles of one-qubit, and display a desirable amount of depth depending on low-dimensional input, however the dimension grows linearly with the number of features, i.e. datasets in high dimensions cause qubit wastage. Feature map encodings, a type of parameterized circuit, which map classical inputs to feature spaces in the Hilbert space are the focus of quantum machine learning, since they directly define the quantity of expressiveness and feature space kernel properties; they become cost-controlled by the depth and parameterization of the feature map circuit. Lastly, QRAM-style architectures are motivated to offer hardware primitives that can permit coherent and parallel access to Classical data in superposition offering sublinear loading cost to an algorithm should they be achievable with desirable latency and fidelity. Nonetheless, architectural studies and reviews indicate that practical cost and other resource arguments against QRAM, particularly in active designs, may be significant and that assertions of hypothetical benefit have to be reconsidered on real values of QRAM conditions. The noted wide taxonomy makes orientation to the analysis of the costs that are to follow it; the encoding decision must reflect the combined limiting demands of data structure, target algorithm, and underlying qubit and control architecture.

Table 1 — Comparative summary of quantum data encoding methods

Encoding Method	Qubit Cost (scaling)	Gate/Depth Characteristic	Typical Advantages	Typical Limitations
Basis encoding	$O(N)$ qubits for N-dimensional one-hot representation	Very shallow; local single-qubit initialisations	Trivial preparing for sparse/indexed data; low depth	Prohibitive qubit count for dense, high-dimensional data
Amplitude encoding	$O(\log N)$ qubits for N-	Gate count often $O(N)$ in naive	High compression of vector	Exact preparation often expensive;

	dimensional vectors	decompositions ; depth depends on factorization	information into few qubits	sensitive to gate errors
Angle (rotation) encoding	$O(d)$ qubits for d features	Shallow per-feature rotations; depth increases with entangling layers	Simple to implement; interpretable per-feature mapping	Scales linearly with features; limited entangling expressivity if shallow
Feature-map (parameterise d circuit)	$O(n_q)$ where n_q is qubits used by circuit	Depth depends on chosen ansatz; can be deep for expressive maps	Tunable expressivity; central to QML kernels	May be expensive to evaluate/classically interpret; design challenge
QRAM-based loading	$O(N)$ classical memory; $O(\log N)$ logical address qubits	Query latency depends on QRAM topology; coherence across routing nodes required	Potential sublinear query cost in idealized models	Hardware complexity, crosstalk, fidelity, and opportunity-cost concerns
MPS-based encoding	$O(k \log N)$ qubits for bond-dimension-limited MPS	Shallow circuits when representing low-entanglement states	Efficient for functions/distributions with low entanglement	Not universally applicable; requires structure in data/function
Polynomial amplitude encoding	$O(\log N)$ qubits	Shallow circuits for low-degree polynomials; controllable approximation	Efficient for smooth/structured functions; reduces gate counts	Approximation error for general functions; limited to certain classes
Sparse-approximation encoding	$O(\log N)$ qubits (if sparsity exploited)	Gate counts scale with number of significant components	Greatly reduces cost for sparse inputs	Requires effective classical sparsification preprocessing
Block-encoding / embedding	$O(n_q + \text{ancilla})$ qubits	Circuit depth depends on block size and controlled unitaries	Facilitates implementation of operators and linear algebra subroutines	Ancilla overhead and nontrivial compilation cost
Hybrid classical-quantum pipelined loading	Varies; often modest qubit counts	Lower quantum depth, higher classical preprocessing	Practical for NISQ-era deployments; reduces quantum burden	Offloads complexity to classical systems; limits end-to-end quantum speedup

4. precision versus imprecision in state preparation Trade-offs and algorithms.

Precise state preparation provides an accurate modeling of the classical data in the desired quantum state; but the level of precision can be prohibitive. Approximate state preparation is the understanding that most algorithms, especially machine learning algorithms and some approximate numerical solvers are resistant to known errors on input encoding. Recent improvements in algorithms produce principled approximation schemes that control the total error in approximation caused by inaccurate amplitude coefficients, but experience a tremendous reduction in the number of gates. Entropy-reduction-based approximate amplitude encoding, truncated singular value decompositions in the form of low-rank quantum circuits and representation-based methods such as the matrix product state (MPS) encodings are some of such methods which use the low entanglement structure in many function classes. The cases of the approximate polynomial amplitude system of discretised polynomials give a specific example where one can readily prepare states representing polynomials to high accuracy using both shallow circuits and controlled ancilla, which allows quantum subroutines to differential equations or spectral methods to be executed without paying the worst-case encodings overhead. Such approximation models usually demand budgetary attention: how much of the allowed error in the algorithm is used in his input encoding versus intrinsic operations (e.g. oracle evaluations, variational optimisations). Compiler-based partitioning of error budgets and resource-conscious tradeoff is a current topic of study and recent frameworks have been proposed through which automated structures produce state preparation circuits based on given error tolerances and available hardware resources.

5. QRAM architectures, feasibility, and critiques

An interesting concept is the so-called quantum random access memory (QRAM): a model in which polylogarithmic overheads are eliminated by loading large size classical datasets into a superposed addressable memory would enable many quantum algorithms to gain practical power. A number of hardware proposals in QRAM such as bucket-brigade, hybrid designs and tree based routing have been made to supplement theoretical algorithm proposals. However, recent surveys and critical evaluations herein point out that QRAM is a fallopian and realistic implementations involve non trivial resource costs that start reintroducing classical comparators. As an illustration, designs of active QRAMs have had arguments over opportunity cost claiming that the control systems and qubit resources used to coherently route data could be instead used to execute very parallel classical precomputation or local classical calculations that reduce the net benefit. Additionally, QRAM emits fidelity costs: with an increase in the addressable memory, the relationship between the routing node and reduction of crosstalk and error

propagation becomes harder and harder, and these error budgets also interact with the error tolerance of the algorithms the addressable memory supports [5,8,51-52]. Recent architecture Research is in hybrid and fat-tree QRAM topologies, which seek to pipeline queries and minimise latencies, and progress is being made towards circuit-level QRAM implementations with limited T-count and depth trade-offs; it remains an end of life cost, but these implementations current architectural studies must ensure that QRAM provides end-to-end speedup costs are properly assessed. In this way, the question of whether QRAM can be loaded with data having a large scale and at a low cost is an open one depending on physical engineering discoveries, and also on favourable behaviour of the algorithm to QRAM-specific errors.

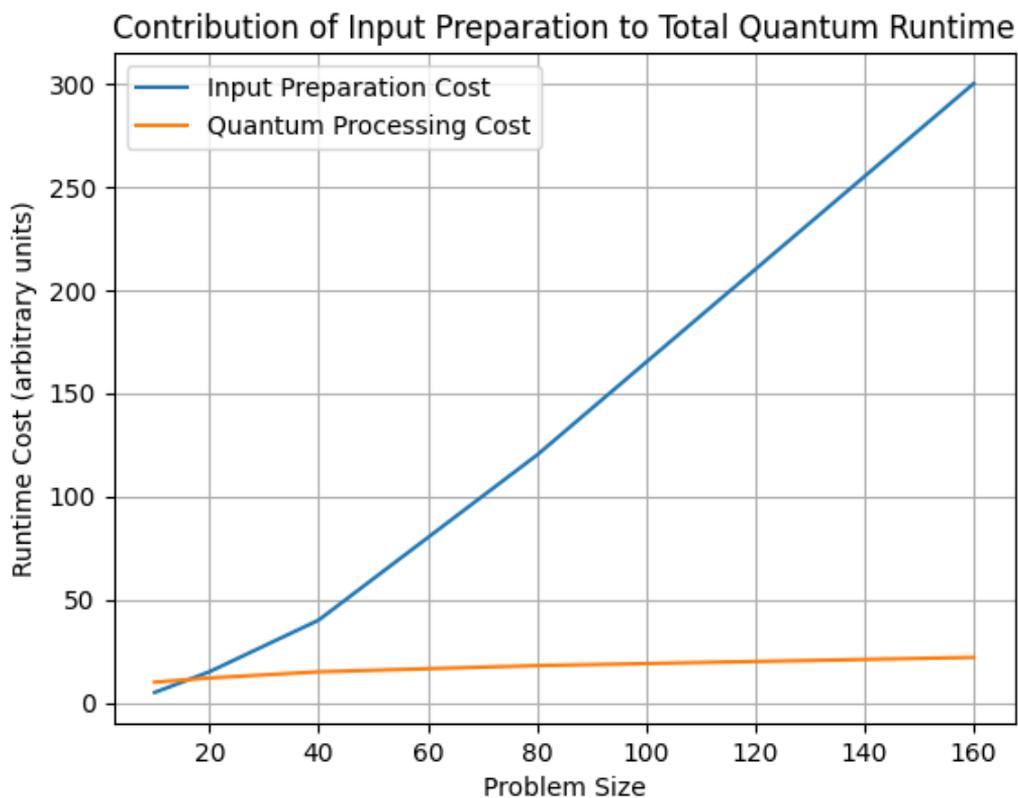


Fig 2: End-to-End Runtime Contribution of Input Preparation

6. Compiler and algorithmic techniques lowering costs of inputs.

Since the worst-case cost of preparing arbitrarily the state of the machine is so unfavorable, a significant body of work is devoted to compiler and algorithmic techniques that influence the manner in which classical data is represented, as well as the quantum circuits which compute it. These methods can be in many formations:

classical preprocessing can compress or sparsify data; analytical transformations can reveal the presence of a polynomial structure and low-entanglement structure; loading only the most significant parts of a vector can be made through the use of basis rotations and partial amplitude encodings; hybrid classical-quantum pipeline can isolate the computationally-beneficial subroutines and leave the heavy-lifting to the classical system, which is going to format data. Techniques on the quantum end, decomposition using Givens rotations, recursive state factorization, and ancilla and multiplexed rotations can be used to lower the number of gates in cases where vectors have structure exploitable by the algorithm, or where approximation errors are tolerated by an algorithm. Recent suggestions on automatic compilation structures explicitly consider data loading a problem in compilation: preserving a target vector, an error budget, and limiting hardware, these compilers generate circuits to prepare the state that optimize circuit count, depth or faithfulness. This treatise stresses that the self-scaling to realistic dataset sizes will be made successful through a tight integration of classical compression, representational assumption and quantum compilation a fact that is tending to find an echo in the literature on co-designing hardware and software.

7. Implications of noise, error accumulation and fault-tolerance.

The cost of input preparation cannot be thought of individually since errors that are added during preparation spread to the remainder of the algorithm, which in many cases are not always nontrivial. Coherent errors in rotations and entangling gates can cause prepared amplitudes to be biased and downstream estimations to be spoiled in devices based on noisy intermediate-scale quantum (NISQ). One can use error mitigation techniques which may need extra queries or resources that again inflate the success cost. In full fault-tolerant designs the cost of state preparation remains a substantial variable: The T-count and T-depth of T-circuits to load data counts are relevant since T gates are costly in most error-correction codes; optimization of T-count in state preparation or QRAM-loading T-circuits can consequently reduce the logic execution time and costs of magic-state distillation overheads by a significant factor. Certain application points of QRAM propose using Clifford+T implementations in a way that minimizes T-depth and T-count, with that minimization being in recognition of the fact that at the fault-tolerant level, circuit-level resource accounting makes the choice of encoding strategy practicable. Below therefore, the credibility of a scaling analysis of preparing inputs needs to consider noise models, error-correction overheads and interaction between the accuracy of encoding and numerical stability of the algorithm. This provides incentives to the use of approximate encodings that purposely give up T-cost to achieve a limited approximation error and methods that allocate error budgets between preparation, computation and measurement phases.

Table 2 — Representative state preparation protocols and resource characteristics

Protocol / Method	Asymptotic Complexity (gates/depth)	Approximation / Error Tradeoff	Hardware Assumptions
Recursive Givens rotations (exact)	$O(N)$ single- and two-qubit rotations; depth $O(N)$ in naive implementations	Exact if implemented without approximation	Requires reliable multi-qubit gates and ancilla control
Bucket-brigade QRAM query	Query-time polylogarithmic in N under ideal model	Error accumulates across routing; needs error correction	Coherent routing nodes with low crosstalk; scalable control
Approximate amplitude encoding (entropy reduction / ER-AAE)	Lower gate count than exact; depends on entropy of vector	Controlled approximation by pruning small amplitudes	Classical preprocessing to identify significance
Polynomial amplitude encoding	Gate count $\text{poly}(\log N, \text{degree})$ for structured polynomials	Approximation error controlled by polynomial degree and discretization	Efficient when target function is low-degree; uses ancilla
Matrix product state (MPS) preparation	Gate depth scales with bond dimension and $\log N$	Approximation error tied to truncation of bond spectrum	Effective when data/function has low entanglement
ER-AAE / entropy-reduction AAE	Sub-linear gate reductions for compressible vectors	Approximates distribution by reducing entropy support	Requires classical sorting/thresholding
Givens + multiplexed rotation compilation	$O(\text{number of non-zero components})$ gates	Exact if full decomposition used; approximate via pruning	Efficient for sparse or structured inputs
Clifford+T optimized QRAM circuits	T-count and T-depth minimized relative to naive implementations	Approximation via bounded T-resource tradeoffs	Fault-tolerant architecture with magic-state distillation
Truncated SVD / low-rank preparation	Gates scale with retained rank r and $\log N$	Error bounded by discarded singular values	Effective when data has low effective rank
Hybrid loading + local rotations	Gate counts modest; classical preprocessing increases	Approximation depends on classical compression fidelity	Practical for near-term devices without QRAM

8. Real life examples: machine learning, linear systems and quantum simulations.

As a way of making these abstract costs tangible, one can take three case studies quantum machine learning (QML), quantum linear algebra solvers and quantum simulation. The expressivity of quantum models in QML has a notion of feature maps; feature maps that are overly expressive can be costly to implement and report and simpler, though also easier, to manage feature maps [9,53]. The choice of data encoding will define the quantized information on the kernels of quantum kernel implemented schemes and determine the sample complexity and generalisation behaviour, empirical data indicate that quantum embedding can ameliorate the classification measurements in certain tasks, and that performance highly depends on data distribution and architecture of the encoding encode. In the case of banded systems solvers such as HHL, imagining an amplitude-encoded right hand side is required, and as long as the state preparation requires linear time in N to be true, the end to end complexity of such an algorithm may lose the benefit of theory. In problems of quantum simulation, structured encodings (e.g. the representation of functions as low-order polynomials or MPS representations) can result in (notably less expensive) shallow circuit models of appropriate physical observables. In all these illustrations, it is the same theme: extension to structure deserve, and encodings, which are approximation-tolerant and structure-aware, can afford practical advantages; whilst naive, worst, encodings can refute the theoretical advantages. Substantiating the importance of choices of data encoding technique in the feasibility of near-term quantum applications, both empirical and systematic reviews of encoding methods support the effectiveness of the encoding choice.

9. Conclusion

The major component that will dictate the future is quantum data encoding which will decide whether quantum algorithms will provide practical benefit or not. The chapter has contended that basic assumptions regarding the data-loading can be negligible and deserve close consideration and attention and that the constraints of representations, approximation, and hardware are to be taken into account. Some of the promising directions are those that support representation-conscious encodings (e.g., polynomial encodings and MPS encodings), automated and resource-conscious compilers based on tradeoffs between precision and depth and T-cost, and pragmatic studies of the QRAM, expressing such tradeoffs in a transparent way. The way forward will involve cross-disciplinary teams consisting of theoretical algorithmics, compiler technology and hardware engineering, realistic benchmarking and mutually extensible toolchains are needed to spearhead objective comparisons. Although a complete implementation of low-cost and large-scale coherent QRAM is an engineering challenge, the near-term strategy to viable quantum data processing is more straightforward: use the structure

where it can, make reasonable trade-offs where it must, and co-architecture encoding schemes with the game courtesy of the hardware upon which they will execute. The recent empirical and theoretical literature supports this methodology since it has been established that significant decreases in preparation cost can be realized when applying this method to structured data and that these decreases do substantially increase the practicality of quantum-enhanced implementations.

References

- [1] Abdelwahab HR, Rauf A, Chen D. Business students' perceptions of Dutch higher educational institutions in preparing them for artificial intelligence work environments. *Industry and Higher Education*. 2023 Feb;37(1):22-34.
- [2] Abrokwah-Larbi K, Awuku-Larbi Y. The impact of artificial intelligence in marketing on the performance of business organizations: evidence from SMEs in an emerging economy. *Journal of Entrepreneurship in Emerging Economies*. 2024 Jun 13;16(4):1090-117.
- [3] Agarwal P, Swami S, Malhotra SK. Artificial intelligence adoption in the post COVID-19 new-normal and role of smart technologies in transforming business: a review. *Journal of Science and Technology Policy Management*. 2024 Apr 18;15(3):506-29.
- [4] Ahmed AA, Agarwal S, Kurniawan IG, Anantadjaya SP, Krishnan C. Business boosting through sentiment analysis using Artificial Intelligence approach. *International Journal of System Assurance Engineering and Management*. 2022 Mar;13(Suppl 1):699-709.
- [5] Alam A, Mohanty A. Business models, business strategies, and innovations in EdTech companies: integration of learning analytics and artificial intelligence in higher education. In 2022 IEEE 6th Conference on Information and Communication Technology (CICT) 2022 Nov 18 (pp. 1-6). IEEE.
- [6] Alawadhi SA, Zowayed A, Abdulla H, Khder MA, Ali BJ. Impact of artificial intelligence on information security in business. In 2022 ASU International Conference in Emerging Technologies for Sustainability and Intelligent Systems (ICETISIS) 2022 Jun 22 (pp. 437-442). IEEE.
- [7] Badghish S, Soomro YA. Artificial intelligence adoption by SMEs to achieve sustainable business performance: application of technology–organization–environment framework. *Sustainability*. 2024 Feb 24;16(5):1864.
- [8] Beheshti A, Yang J, Sheng QZ, Benatallah B, Casati F, Dustdar S, Nezhad HR, Zhang X, Xue S. ProcessGPT: transforming business process management with generative artificial intelligence. In 2023 IEEE international conference on web services (ICWS) 2023 Jul 2 (pp. 731-739). IEEE.
- [9] Bickley SJ, Macintyre A, Torgler B. Artificial intelligence and big data in sustainable entrepreneurship. *Journal of Economic Surveys*. 2025 Feb;39(1):103-45.
- [10] Bolton C, Machová V, Kovacova M, Valaskova K. The power of human–machine collaboration: Artificial intelligence, business automation, and the smart economy. *Economics, Management, and Financial Markets*. 2018 Dec 1;13(4):51-6.
- [11] Canhoto AI, Clear F. Artificial intelligence and machine learning as business tools: A framework for diagnosing value destruction potential. *Business Horizons*. 2020 Mar 1;63(2):183-93.

- [12] Cavazza A, Dal Mas F, Paoloni P, Manzo M. Artificial intelligence and new business models in agriculture: a structured literature review and future research agenda. *British Food Journal*. 2023 Jul 12;125(13):436-61.
- [13] Chen R, Zhang T. Artificial intelligence applications implication for ESG performance: can digital transformation of enterprises promote sustainable development?. *Chinese Management Studies*. 2025 May 13;19(3):676-701.
- [14] Chowdhury S, Budhwar P, Wood G. Generative artificial intelligence in business: towards a strategic human resource management framework. *British Journal of Management*. 2024 Oct;35(4):1680-91.
- [15] Davenport TH, Ronanki R. Artificial intelligence for the real world. *Harvard business review*. 2018 Jan 1;96(1):108-16.
- [16] Demaidi MN. Artificial intelligence national strategy in a developing country. *Ai & Society*. 2025 Feb;40(2):423-35.
- [17] Drydakakis N. Artificial Intelligence and reduced SMEs' business risks. A dynamic capabilities analysis during the COVID-19 pandemic. *Information Systems Frontiers*. 2022 Aug;24(4):1223-47.
- [18] Fallahi S, Mellquist AC, Mogren O, Listo Zec E, Alguren P, Hallquist L. Financing solutions for circular business models: Exploring the role of business ecosystems and artificial intelligence. *Business Strategy and the Environment*. 2023 Sep;32(6):3233-48.
- [19] Feuerriegel S, Shrestha YR, von Krogh G, Zhang C. Bringing artificial intelligence to business management. *Nature Machine Intelligence*. 2022 Jul;4(7):611-3.
- [20] Gadhave RT, Dhingra SK, Abhishek MB, Thota MK, Sholapurapu PK, Lamba V, Patil AK, Yadav MS. Deep Learning-Enabled Decision Support Systems For Strategic Business Management. *International Journal of Environmental Sciences*. 2025;11(7):2025.
- [21] Gursoy D, Cai R. Artificial intelligence: an overview of research trends and future directions. *International Journal of Contemporary Hospitality Management*. 2025 Jan 2;37(1):1-7.
- [22] Hu KH, Chen FH, Hsu MF, Tzeng GH. Governance of artificial intelligence applications in a business audit via a fusion fuzzy multiple rule-based decision-making model. *Financial Innovation*. 2023 Aug 14;9(1):117.
- [23] Jain S, Sholapurapu PK, Sharma B, Nagar M, Bhatt N, Swaroopa N. Hybrid Encryption Approach for Securing Educational Data Using Attribute-Based Methods. In: 2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 2025 Apr 9 (pp. 1-6). IEEE.
- [24] Jorzik P, Antonio JL, Kanbach DK, Kallmuenzer A, Kraus S. Sowing the seeds for sustainability: A business model innovation perspective on artificial intelligence in green technology startups. *Technological forecasting and social change*. 2024 Nov 1;208:123653.
- [25] Jorzik P, Yigit A, Kanbach DK, Kraus S, Dabić M. Artificial intelligence-enabled business model innovation: Competencies and roles of top management. *IEEE transactions on engineering management*. 2023 May 24;71:7044-56.
- [26] Kalogiannidis S, Kalfas D, Papaevangelou O, Giannarakis G, Chatzitheodoridis F. The role of artificial intelligence technology in predictive risk assessment for business continuity: A case study of Greece. *Risks*. 2024 Jan 23;12(2):19.
- [27] Kanbach DK, Heiduk L, Bluecher G, Schreiter M, Lahmann A. The GenAI is out of the bottle: generative artificial intelligence from a business model innovation perspective. *Review of Managerial Science*. 2024 Apr;18(4):1189-220.

- [28] Kar AK, Kushwaha AK. Facilitators and barriers of artificial intelligence adoption in business—insights from opinions using big data analytics. *Information Systems Frontiers*. 2023 Aug;25(4):1351-74.
- [29] Kasireddy LC, Bhupathi HP, Shrivastava R, Sholapurapu PK, Bhatt N. Intelligent Feature Selection Model using Artificial Neural Networks for Independent Cyberattack Classification. In 2025 2nd International Conference On Multidisciplinary Research and Innovations in Engineering (MRIE) 2025 Jul 30 (pp. 572-576). IEEE.
- [30] Krishnan C, Gupta A, Gupta A, Singh G. Impact of artificial intelligence-based chatbots on customer engagement and business growth. In *Deep learning for social media data analytics 2022* Sep 19 (pp. 195-210). Cham: Springer International Publishing.
- [31] Kulkov I. Next-generation business models for artificial intelligence start-ups in the healthcare industry. *International Journal of Entrepreneurial Behavior & Research*. 2023 May 4;29(4):860-85.
- [32] Kumar S, Lim WM, Sivarajah U, Kaur J. Artificial intelligence and blockchain integration in business: trends from a bibliometric-content analysis. *Information systems frontiers*. 2023 Apr;25(2):871-96.
- [33] Mishra AN, Pani AK. Business value appropriation roadmap for artificial intelligence. *VINE Journal of Information and Knowledge Management Systems*. 2021 May 31;51(3):353-68.
- [34] Nagar M, Sholapurapu PK, Kaur DP, Lathigara A, Amulya D, Panda RS. A Hybrid Machine Learning Framework for Cognitive Load Detection Using Single Lead EEG, CiSSA and Nature-Inspired Feature Selection. In 2025 World Skills Conference on Universal Data Analytics and Sciences (WorldSUAS) 2025 Aug 22 (pp. 1-6). IEEE.
- [35] Paramesha M, Rane N, Rane J. Big data analytics, artificial intelligence, machine learning, internet of things, and blockchain for enhanced business intelligence. *Artificial Intelligence, Machine Learning, Internet of Things, and Blockchain for Enhanced Business Intelligence* (June 6, 2024). 2024 Jun 6.
- [36] Paramesha M, Rane N, Rane J. Big data analytics, artificial intelligence, machine learning, internet of things, and blockchain for enhanced business intelligence. *Artificial Intelligence, Machine Learning, Internet of Things, and Blockchain for Enhanced Business Intelligence* (June 6, 2024). 2024 Jun 6.
- [37] Qi B, Shen Y, Xu T. An artificial-intelligence-enabled sustainable supply chain model for B2C E-commerce business in the international trade. *Technological forecasting and social change*. 2023 Jun 1;191:122491.
- [38] Rana NP, Chatterjee S, Dwivedi YK, Akter S. Understanding dark side of artificial intelligence (AI) integrated business analytics: assessing firm's operational inefficiency and competitiveness. *European Journal of Information Systems*. 2022 May 4;31(3):364-87.
- [39] Rane NL, Paramesha M, Choudhary SP, Rane J. Artificial intelligence, machine learning, and deep learning for advanced business strategies: a review. *Partners Universal International Innovation Journal*. 2024 Jun 25;2(3):147-71.
- [40] Reddy MU, Bhagyalakshmi L, Sholapurapu PK, Lathigara A, Singh AK, Nidadavolu V. Optimizing Scheduling Problems in Cloud Computing Using a Multi-Objective Improved Genetic Algorithm. In 2025 2nd International Conference On Multidisciplinary Research and Innovations in Engineering (MRIE) 2025 Jul 30 (pp. 635-640). IEEE.

- [41] Reier Forradellas RF, Garay Gallastegui LM. Digital transformation and artificial intelligence applied to business: Legal regulations, economic impact and perspective. *Laws*. 2021 Aug 27;10(3):70.
- [42] Sahoo S, Kumar S, Donthu N, Singh AK. Artificial intelligence capabilities, open innovation, and business performance—Empirical insights from multinational B2B companies. *Industrial marketing management*. 2024 Feb 1;117:28-41.
- [43] Sandeep SR, Ahamad S, Saxena D, Srivastava K, Jaiswal S, Bora A. To understand the relationship between Machine learning and Artificial intelligence in large and diversified business organisations. *Materials Today: Proceedings*. 2022 Jan 1;56:2082-6.
- [44] Shaik AS, Alshibani SM, Jain G, Gupta B, Mehrotra A. Artificial intelligence (AI)-driven strategic business model innovations in small-and medium-sized enterprises. Insights on technological and strategic enablers for carbon neutral businesses. *Business Strategy and the Environment*. 2024 May;33(4):2731-51.
- [45] Sholapurapu PK, Omkar J, Bansal S, Gandhi T, Tanna P, Kalpana G. Secure Communication in Wireless Sensor Networks Using Cuckoo Hash-Based Multi-Factor Authentication. In *2025 World Skills Conference on Universal Data Analytics and Sciences (WorldSUAS) 2025 Aug 22* (pp. 1-6). IEEE.
- [46] Shwedeh F, Alzoubi HM. Creating and Evaluating Instructional Java Programming Codes with Utilization of Artificial Intelligence for Customized Business Requirements. In *International Scientific Conference Management and Engineering 2024 Jun 23* (pp. 281-286). Cham: Springer Nature Switzerland.
- [47] Sjödin D, Parida V, Kohtamäki M. Artificial intelligence enabling circular business model innovation in digital servitization: Conceptualizing dynamic capabilities, AI capacities, business models and effects. *Technological Forecasting and Social Change*. 2023 Dec 1;197:122903.
- [48] Toniolo K, Masiero E, Massaro M, Bagnoli C. Sustainable business models and artificial intelligence: Opportunities and challenges. *Knowledge, people, and digital transformation: Approaches for a sustainable future*. 2020 Apr 23:103-17.
- [49] Upadhyay N, Upadhyay S, Al-Debei MM, Baabdullah AM, Dwivedi YK. The influence of digital entrepreneurship and entrepreneurial orientation on intention of family businesses to adopt artificial intelligence: examining the mediating role of business innovativeness. *International Journal of Entrepreneurial Behavior & Research*. 2023 Jan 17;29(1):80-115.
- [50] Vardarlier P, Zafer C. Use of artificial intelligence as business strategy in recruitment process and social perspective. In *Digital business strategies in blockchain ecosystems: Transformational design and future of global business 2019 Nov 10* (pp. 355-373). Cham: Springer International Publishing.
- [51] William P, Panicker A, Falah A, Hussain A, Shrivastava A, Khan AK. The Emergence of Artificial Intelligence and Machine Learning in Contemporary Business Management. In *2023 4th International Conference on Computation, Automation and Knowledge Management (ICCAKM) 2023 Dec 12* (pp. 1-6). IEEE.
- [52] Zhang X, Wang L, Helwig J, Luo Y, Fu C, Xie Y, Liu M, Lin Y, Xu Z, Yan K, Adams K. Artificial intelligence for science in quantum, atomistic, and continuum systems. *Foundations and Trends® in Machine Learning*. 2025 Jul 20;18(4):385-912.
- [53] Zhou X, Li G, Wang Q, Li Y, Zhou D. Artificial intelligence, corporate information governance and ESG performance: Quasi-experimental evidence from China. *International Review of Financial Analysis*. 2025 Jun 1;102:104087.