QTML 2021

Quantum Techniques in Machine Learning

SHORT ABSTRACTS

Monday Nov 8 2021

Session S1.1

Invited Talk 9:25
Mon Nov 8

The quest to open the black box of Deep Learning

Sanjeev Arora
Princeton University

Deep learning has rapidly come to dominate AI and machine learning in the past decade. These successes have come despite deep learning largely being a "black box." Now a small sub-discipline has grown up trying to derive better understanding of the underlying mathematical properties. Via a tour d'horizon of recent theoretical analyses of deep learning in some concrete settings, we illustrate how the black box view can miss out on (or even be wrong about) special phenomena going on during training. These phenomena are also not captured by the training objective. We argue that understanding such phenomena via mathematical understanding will be crucial for enabling the full range of future applications.

Paper 77: Quantum Algorithms for Reinforcement Learning with a Generative Model

Daochen Wang, Aarthi Sundaram, Robin Kothari, Ashish Kapoor and Martin Roetteler

Reinforcement learning studies how an agent should interact with an environment to maximize its cumulative reward. A standard way to study this question abstractly is to ask how many samples an agent needs from the environment to learn an optimal policy for a γ-discounted Markov decision process (MDP). For such an MDP, we design quantum algorithms that approximate an optimal policy (π^*), the optimal value function (v^*), and the optimal Q-function (q^*).

Paper 24: Universal Compiling and (No-)Free-Lunch Theorems for Continuous Variable Quantum Learning

Tyler Volkoff, Zoe Holmes and *Andrew Sornborger*

We introduce algorithms for continuous-variable (CV) variational quantum compiling which are motivated by extending "no-free-lunch" theorems to the CV setting. These algorithms utilize Gaussian resources, e.g., coherent states and squeezed states, and involve variational circuits that are efficiently trainable and scalable, avoiding common obstructions such as the barren plateau phenomenon.

Session S1.2

Barren Plateaus and Quantum Generative Training Using Rényi Divergences

Invited Talk 11:30
Mon Nov 8

Nathan Wiebe
University of Toronto

Recently there has been substantial interest in the development of quantum machine learning protocols. However, despite this there have been a number of stumbling blocks that have emerged for quantum machine learning. Specifically, recent work showing dequantizations of quantum algorithms as well as barren plateau results have shown that new ideas are needed to build and train quantum models that have advantage over classical models. In this talk I will discuss the challenges that arise when trying to use quantum mechanical effects to train neural networks and show that entanglement, long viewed as a boon for quantum computers, is actually generically anathema for deep quantum neural networks and by drawing on lessons from the thermalization community that unchecked entanglement causes these models to be exponentially close to the maximally mixed state and that gradient descent is not capable of remedying the situation. I will then look at a partial solution to these problems, which involves switching to generative learning and show that quantum neural networks can be efficiently trained using quantum analogues of the KL-divergence, which do not suffer from barren plateau problems due to its logarithmic divergence for orthogonal states (which combats gradient decay). I will show numerical evidence that indicates that small scale quantum neural networks can be trained to generate complex quantum states and in turn suggest that generatively pre-training models may be a way to circumvent known barren plateau results for quantum machine learning.

Paper 37: Enhancing Combinatorial Optimization with Quantum Generative Models

Francisco Javier Fernandez Alcazar and Alejandro Perdomo Ortiz

Combinatorial optimization is one of the key candidates in the race for practical quantum advantage. In this work, we introduce a new family of quantum-enhanced optimizers and demonstrate how quantum machine learning models known as quantum generative models can find lower minima than those found by means of state-of-the-art classical solvers.

Paper 76: Diagnosing barren plateaus with tools from quantum optimal control

Martin Larocca, Marco Cerezo, Patrick Coles, Kunal Sharma, Piotr Czarnik, and Gopikrishnan Muraleedharan

Variational Quantum Algorithms (VQAs) have received considerable attention due to their potential for achieving near-term quantum advantage. However, the observation of barren plateaus, a phenomenon by which the landscape becomes exponentially flat in the number of qubits, has raised uncertainty around their scalability. In this work we employ tools from quantum optimal control to develop a framework that can diagnose the presence, or absence, of barren plateaus for a certain class of periodic ansatzes.

Paper 74: Adaptive shot allocation for fast convergence in variational quantum algorithms

Andrew Arrasmith, Andi Gu, Angus Lowe, Pavel Dub and Patrick Coles

Variational Quantum Algorithms (VQAs) are a promising approach for practical applications like chemistry and materials science on near-term quantum computers as they typically reduce quantum resource requirements. However, in order to implement VQAs, an efficient classical optimization strategy is required. Here we present a new stochastic gradient descent method using an adaptive number of shots at each step, called the global Coupled Adaptive Number of Shots (gCANS) method, which improves on prior art in both the number of iterations as well as the number of shots required. These improvements reduce both the time and money required to run VQAs on current cloud platforms. We analytically prove that in a convex setting gCANS achieves geometric convergence to the optimum. Further, we numerically investigate the performance of gCANS on some chemical configuration problems. We also consider finding the ground state for an Ising model with different numbers of spins to examine the scaling of the method. We find that for these problems, gCANS compares favorably to all of the other optimizers we consider.

Session S1.3

Quantum machine learning: quantum kernel method and beyond

Invited Talk 14:30
Mon Nov 8

Keisuke Fujii
Osaka University

In this talk, I will present the latest results on quantum machine learning as a variational quantum algorithm expected to be effective on NISQ devices. In particular, we will provide an approach that is essentially different from the kernel method, which is known to be able to characterize quantum machine learning, including our proposed quantum circuit learning. This new approach provides a better performance than the quantum kernel method. Interestingly, this new approach is also characterized by another kernel known as quantum "tangent" kernel in the limit of over parameterization. Regarding an application, we will also provide an application of the quantum kernel like clustering method to classification of quantum phases in quantum many-body systems, as a task in which quantum machine learning may have an advantage.

Paper 3: Exploring Quantum Perceptron and Quantum Neural Network structures with a teacher-student scheme

Aikaterini Gratsea and *Patrick Huembeli*

Near-term quantum devices can be used to build quantum machine learning models, such as quantum kernel methods and quantum neural networks (QNN) to perform classification tasks. There have been many proposals on how to use variational quantum circuits as quantum perceptrons or as QNNs. The aim of this work is to introduce a teacher-student scheme that could systematically compare any QNN architectures and evaluate their relative expressive power. Specifically, the teacher model generates the datasets mapping random inputs to outputs which then have to be learned by the student models. This way, we avoid training on arbitrary data sets and compare the learning capacity of different models directly via the loss, the prediction map, the accuracy and the relative entropy between the prediction maps. Here, we focus particularly on a quantum perceptron model inspired by the recent work of Tacchino et. al. [1] and compare it to the data re-uploading scheme that was originally introduced by Pérez-Salinas et. al. [2]. We discuss alterations of the perceptron model and the formation of deep QNN to better understand the role of hidden units and the non-linearities in these architectures.

[1] F. Tacchino, C. Macchiavello, D. Gerace, and D. Bajoni, npj Quantum Information 5(2019), 10.1038/s41534-019-0140-4.

[2] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I.Latorre, Quantum4, 226 (2020).

Paper 50: Towards understanding the power of quantum kernels in the NISQ era

Xinbiao Wang, Yuxuan Du, Yong Luo, and Dacheng Tao

A key problem in the field of quantum computing is understanding whether quantum machine learning (QML) models implemented on noisy intermediate-scale quantum (NISQ) machines can achieve quantum advantages. Recently, Huang et al. [Nat Commun 12, 2631] partially answered this question by the lens of quantum kernel learning. Namely, they exhibited that quantum kernels can learn specific datasets with lower generalization error over the optimal classical kernel methods. However, most of their results are established on the ideal setting and ignore the caveats of near-term quantum machines. To this end, a crucial open question is: does the power of quantum kernels still hold under the NISQ setting? In this study, we fill this knowledge gap by exploiting the power of quantum kernels when the quantum system noise and sample error are considered. Concretely, we first prove that the advantage of quantum kernels vanishes for large size of datasets, few number of measurements, and large system noise. With the aim of preserving the superiority of quantum kernels in the NISQ era, we further devise an effective method via indefinite kernel learning. Numerical simulations accord with our theoretical results. Our work provides theoretical guidance of exploring advanced quantum kernels to attain quantum advantages on NISQ devices.

Paper 47: Quantum evolution kernel : Machine learning on graphs with programmable arrays of qubits

Slimane Thabet, Louis-Paul Henry, Constantin Dalyac, and Loic Henriet

The rapid development of reliable Quantum Processing Units (QPU) opens up novel computational opportunities for machine learning. Here, we introduce a procedure for measuring the similarity between graph-structured data, based on the time-evolution of a quantum system. By encoding the topology of the input graph in the Hamiltonian of the system, the evolution produces measurement samples that retain key features of the data. We study analytically the procedure and illustrate its versatility in providing links to standard classical approaches. We then show numerically that this scheme performs well compared to standard graph kernels on typical benchmark datasets. Finally, we study the possibility of a concrete implementation on a realistic neutral-atom quantum processor.

Session S1.4

Hintons, Quantum Fields and Deep Learning

Invited Talk 17:00
Mon Nov 8

Max Welling
University of Amsterdam and Microsoft Research

While it is highly speculative whether quantum computation is going to help machine learning, quantum mechanics, and in particular quantum field theory (QFT), does provide a unique language to represent deep learning. In this talk we will explore how to use the language of QFT to represent a neural network, and muse about ways in which this could be useful for computation. In particular, the QFT representation seems well suited for optical (quantum) computation, but we make no claims as to whether this means deep learning will really benefit from quantum computers.

Paper 82: Testing identity of collections of quantum states: sample complexity analysis

Marco Fanizza, Raffaele Salvia, and Vittorio Giovannetti

We study the problem of testing identity of a collection of unknown quantum states given sample access to this collection, each state appearing with some known probability. We show that for a collection of d-dimensional quantum states of cardinality N, the sample complexity is $O(\sqrt{N}d/\epsilon^2)$, which is optimal up to a constant. The test is obtained by estimating the mean squared Hilbert-Schmidt distance between the states, thanks to a suitable generalization of the estimator of the Hilbert-Schmidt distance between two unknown states by Bădescu, O'Donnell, and Wright (https://dl.acm.org/doi/10.1145/3313276.3316344).

Paper 61: Complexity of Quantum Support Vector Machines and Quantum Neural Networks

Arne Thomsen, David Sutter, Amira Abbas, and Stefan Woerner

We prove a polynomial speedup compared to [Liu, et al. "A rigorous and robust quantum speed-up in supervised machine learning"] for training of noisy quantum support vector machines via the dual optimization problem. We introduce the Pegasos algorithm [Shalev-Shwartz, et al. "Pegasos: Primal Estimated sub-GrAdient SOlver for SVM"] as an alternative and derive bounds on its runtime, which scales favorably. In addition, we analyze quantum neural networks numerically from the same perspective.

Paper 6: Linear Regression by Quantum Amplitude Estimation and its Extension to Convex Optimization

Kazuya Kaneko, *Koichi Miyamoto*, Naoyuki Takeda, and Kazuyoshi Yoshino

We propose a quantum algorithm for linear regression based on quantum amplitude estimation. This outputs regression coefficients as classical data with complexity depending on the number of data logarithmically and on the tolerance ϵ as $O(\epsilon^{-1})$, in contrast to $O(\epsilon^{-2})$ in existing methods. Additionally, we generalize it for convex optimization.

Tuesday Nov 9 2021

Session S2.1

Machine learning for quantum information: learning to generate quantum optical setups and quantum dynamics

Invited Talk 9:00
Tue Nov 9

Alan Aspuru-Guzik
University of Toronto

In this talk I will talk about two recent works from my laboratory. I will begin by discussing the development of a generative approach that generates quantum optical setups for generating entangled quantum states. The latent space of such autoencoder is a powerful tool for the exploration and interpretation of novel experimental setups. In the second part of the talk, I will discuss our work to predict decoherence dynamics in idealized qubits using time-series machine learning methods.

Paper 31: FLIP: A flexible initializer for arbitrarily-sized parametrized quantum circuits

Frederic Sauvage, Alejandro Perdomo-Ortiz, Sukin Sim, Alexander Kunitsa, William Simon, and Marta Mauri

Developing efficient methods for the training of quantum circuits is critical to the success of any variational quantum algorithms. We address this task from an initialization perspective and propose a novel meta-learning scheme which not only surpasses state-of-the-art strategies but also allows to initialize larger circuits than used during training.

Paper 70: Impact of Noise and Error Mitigation on Trainability of Variational Quantum Algorithms

Samson Wang, Piotr Czarnik, Andrew Arrasmith, Marco Cerezo, Enrico Fontana, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick Coles

Variational Quantum Algorithms (VQAs) are widely viewed as the best hope for near-term quantum advantage. In the first part of this talk, we show that noise can severely limit the trainability of VQAs by exponentially flattening the cost landscape and suppressing the magnitudes of cost gradients. Error Mitigation (EM) shows promise in reducing the impact of noise on near-term devices. Thus, it is also natural to ask whether EM can improve the trainability of VQAs. In the second part of this talk, we show that, for a broad class of EM strategies, exponential cost concentration cannot be resolved without committing exponential resources elsewhere. This class of strategies includes as special cases Zero Noise Extrapolation, Virtual Distillation, Probabilistic Error Cancellation, and Clifford Data Regression. In addition, we perform analytical and numerical analysis of these

EM protocols, and we find that some of them (e.g., Virtual Distillation) can make it harder to resolve cost function values compared to running no EM at all. As a positive result, we do find numerical evidence that Clifford Data Regression (CDR) can aid the training process in certain settings where cost concentration is not too severe. Our results show that care should be taken in applying EM protocols as they can either worsen or not improve trainability. On the other hand, our positive results for CDR highlight the possibility of engineering error mitigation methods to improve trainability.

Paper 49: Scalable Bosonic Random Walk Networks for Graph Learning

Shivshankar and Don Towsley

The development of Graph Neural Networks (GNNs) has led to great progress in machine learning on graph-structured data. These networks operate via diffusing information across the graph nodes while capturing the structure of the graph. Recently there has also seen tremendous progress in quantum computing techniques. In this work, we explore applications of multi-particle quantum walks on diffusing information across graphs. Our model is based on learning the operators that govern the dynamics of quantum random walkers on graphs. We demonstrate the effectiveness of our method on a temperature prediction based regression task.

Session S2.2

Information-Theoretic Bounds on Quantum Advantage in Machine Learning

Invited Talk 11:30
Tue Nov 9

Hsin-Yuan (Robert) Huang
California Institute of Technology

We study the performance of classical and quantum machine learning (ML) models in predicting outcomes of physical experiments. The experiments depend on an input parameter x and involve execution of a (possibly unknown) quantum process E. Our figure of merit is the number of runs of E required to achieve a desired prediction performance. We consider classical ML models that perform a measurement and record the classical outcome after each run of E, and quantum ML models that can access E coherently to acquire quantum data; the classical or quantum data is then used to predict outcomes of future experiments. We prove that for any input distribution D(x), a classical ML model can provide accurate predictions on average by accessing E a number of times comparable to the optimal quantum ML model. In contrast, for achieving accurate prediction on all inputs, we prove that exponential quantum advantage is possible. For example, to predict expectations of all Pauli observables in an n-qubit system ρ, classical ML models require $\Omega(2^{n/3})$ copies of ρ, but we present a quantum ML model using only O(n) copies. Our results clarify where quantum advantage is possible and highlight the potential for classical ML models to address challenging quantum problems in physics and chemistry. The arXiv link: https://arxiv.org/abs/2101.02464.

Paper 81: Generative Quantum Learning of Joint Probability Distribution Functions

Elton Yechao Zhu, Sonika Johri, Dave Bacon, Mert Esencan, Jungsang Kim, Mark Muir, Nikhil Murgai, Jason Nguyen, Neal Pisenti, Adam Schouela, Ksenia Sosnova, and Ken Wright

Modeling joint probability distributions is an important task in a wide variety of fields. One popular technique for this employs a family of multivariate distributions with uniform marginals called copulas. While the theory of modeling joint distributions via copulas is well understood, it gets practically challenging to accurately model real data with many variables. In this work, we design quantum machine learning algorithms to model copulas. We show that any copula can be naturally mapped to a multipartite maximally entangled state. A variational ansatz we christen as a 'qopula' creates arbitrary correlations between variables while maintaining the copula structure starting from a set of Bell pairs for two variables, or GHZ states for multiple variables. As an application, we train a Quantum Generative Adversarial Network (QGAN) and a Quantum Circuit Born Machine (QCBM) using this variational ansatz to generate samples from joint distributions of two variables for historical data from the stock market. We demonstrate our generative learning algorithms on trapped ion quantum computers from IonQ for up to 8 qubits and show

that our results outperform those obtained through equivalent classical generative learning. Further, we present theoretical arguments for exponential advantage in our model's expressivity over classical models based on communication and computational complexity arguments.

Paper 23: Estimating distinguishability measures on quantum computers

Rochisha Agarwal, Soorya Rethinasamy, Kunal Sharma, and Mark M. Wilde

The performance of a quantum information processing protocol is ultimately judged by distinguishability measures that quantify how distinguishable the actual result of the protocol is from the ideal case. The most prominent distinguishability measures are those based on the fidelity and trace distance, due to their physical interpretations. In this paper, we propose and review several algorithms for estimating distinguishability measures based on trace distance and fidelity, and we evaluate their performance using simulators of quantum computers. The algorithms can be used for distinguishing quantum states, channels, and strategies (the last also known in the literature as "quantum combs"). The fidelity-based algorithms offer novel physical interpretations of these distinguishability measures in terms of the maximum probability with which a single prover (or competing provers) can convince a verifier to accept the outcome of an associated computation. We simulate these algorithms by using a variational approach with parameterized quantum circuits and find that they converge well for the examples that we consider.

Session S2.3

Invited Talk 14:30
Tue Nov 9

Experimental realization of quantum machine learning algorithms

Kosuke Mitarai
Osaka University

Recent proposals for NISQ machine learning algorithms necessitate their experimental demonstrations for evaluating their performances. I will discuss the use of ensemble quantum systems for such purposes. Ensemble systems allow us to obtain expectation values of observables with one-shot experiments. This property is useful for NISQ algorithms since many of them mainly use quantum computers to estimate expectation values. In this talk, I will present our recent experimental demonstrations of NISQ machine learning algorithms using nuclear spins in solid controlled with NMR technique. The sophisticated experimental technique of NMR allows us to involve more than 25 spins in the algorithm, which is one of the largest quantum systems employed for machine learning purposes to this day. I will also compare the results with classical machine learning techniques and discuss the limitations.

Paper 38: Storage capacity and learning capability of quantum neural networks

Aikaterini Gratsea, Maciej Lewenstein, Anna Sanpera, Albert Alloy, Andreu Riera-Campeny, and Valentin Kasper

We study the storage capacity of quantum neural networks (QNNs), described by completely positive trace preserving (CPTP) maps acting on a N-dimensional Hilbert space. We demonstrate that attractor QNNs can store in a non-trivial manner up to N linearly independent pure states. For n qubits, QNNs can reach an exponential storage capacity, $\mathcal{O}(2^n)$, clearly outperforming standard classical neural networks whose storage capacity scales linearly with the number of neurons n. We estimate, employing the Gardner program, the relative volume of CPTP maps with $M \leq N$ stationary states and show that this volume decreases exponentially with M and shrinks to zero for $M \geq N + 1$. We generalize our results to QNNs storing mixed states as well as input-output relations for feed-forward QNNs. Our approach opens the path to relate storage properties of QNNs to the quantum features of the input-output states. This paper is dedicated to the memory of Peter Wittek.

Paper 16: Quantum Circuits for Graph Representation Learning

Peter Mernyei, Konstantinos Meichanetzidis, and İsmail İlkan Ceylan

We investigate quantum circuits for learning functions over graphs, proposing a unifying framework of Equivariant Quantum Graph Circuits subsuming earlier work as special cases and discuss other possible subclasses. We prove that these circuits are universal approximators for functions over the graph domain and provide experimental evidence.

Paper 32: Large-scale quantum machine learning

Tobias Haug, Chris N. Self, and Myungshik Kim

We substantially speed up quantum machine learning of large datasets via randomized measurements of quantum kernels. We identify a high-dimensional data encoding that is characterized by quantum geometric measures. Our technique successfully classifies handwritten images using IBM quantum computers with parallel processing and a complementary noise mitigation method.

Session S2.4

Invited Talk 17:00
Tue Nov 9

Why machine learning with quantum computers often reduces to kernel methods

Maria Schuld
Xanadu/University of KwaZulu-Natal

With the race for quantum computers in full swing, researchers became interested in the question of what happens if we replace a machine learning model with a quantum circuit. While such "quantum models" are sometimes called "quantum neural networks", their mathematical structure reveals that they are in fact kernel methods with kernels that measure the distance between data embedded into quantum states. This talk gives an overview of the link, and explains what it means for quantum and classical machine learning.

Paper 57: The Inductive Bias of Quantum Kernels

Jonas Kübler, Simon Buchholz, and Bernhard Schölkopf

We analyze the spectral properties of quantum kernels and show how they determine the potential for a quantum advantage in supervised machine learning. We show a no-free lunch theorem and furthermore that kernel evaluations might block the road to quantum advantages, similarly to Barren Plateaus in QNNs.

Paper 30: Universal Approximation Property via Quantum Feature Maps

Quoc Hoan Tran, Takahiro Goto, and Kohei Nakajima

The quantum Hilbert space can be used as a quantum-enhanced feature space in machine learning (ML) via the quantum feature map to encode classical data into quantum states. We prove the ability to approximate any continuous function with optimal approximation rate via quantum ML models in typical quantum feature maps.

Wednesday Nov 10 2021

Session S3.1

Efficient Quantum Optimization via Multi-Basis Encodings and Tensor Rings

Invited Talk 9:00
Wed Nov 10

Anima Anandkumar
California Institute of Technology

Despite extensive research efforts, few quantum algorithms for classical optimization demonstrate an advantage that is realizable on near-term devices. The utility of many quantum algorithms is limited by high requisite circuit depth and nonconvex optimization landscapes. We tackle these challenges by introducing a new variational quantum algorithm that utilizes multi-basis graph encodings and nonlinear activation functions. Our technique results in increased optimization performance, a factor of two increase in effective quantum resources, and a quadratic reduction in measurement complexity. Further, we construct exact circuit representations using factorized tensor rings. This enables us to successfully optimize the MaxCut of the non-local 512-vertex DIMACS library graphs on a single A100 GPU using only shallow circuits. We further provide efficient distributed implementation via the Tensorly-Quantum library.

Paper 4: Quantum Advantage in Basis-Enhanced Neural Sequence Models

Eric R. Anschuetz and Xun Gao

Quantum mechanical systems are able to produce probability distributions that are difficult to capture using classical models. Using this principle, it was previously shown that certain quantum extensions of Bayesian networks were superpolynomially more expressive than their classical counterparts when applied to a certain sequence-to-sequence learning task. It was unclear whether this advantage would carry over to neural sequence models, which are generally exponentially more efficient than Bayesian networks. Here, we show that due to the Lipschitz continuity of both seq2seq and Transformer based models under reasonable trainability assumptions, a superpolynomial separation remains between neural sequence learning models and their quantum counterparts. Due to the simplicity of the quantum extension, we are also able to directly link the former separation to quantum contextuality and the latter to quantum complementarity. We also show that there is a polynomial expressivity separation between quantum extended sequence models and arbitrary autoregressive sequence models due to quantum contextuality.

Paper 2: Kernel Matrix Completion for Offline Quantum-Enhanced Machine Learning

Annie Naveh, Anna Phan, Imogen Fitzgerald, Andrew Lockwood, and Travis Scholten

Hybrid quantum-classical machine learning workflows involving quantum kernel matrices incur data transfer costs scaling quadratically with the data set size. This work shows matrix completion algorithms mitigate these costs while being robust to shot noise. The relationship between properties of quantum circuits and completability of quantum kernel matrices is studied.

Paper 25: Quantum-inspired manifold learning

Akshat Kumar and *Mohan Sarovar*

We introduce a (classical) algorithm for extracting geodesic distances on sampled manifolds that relies on simulation of quantum dynamics on a graph embedding of the sampled data. Our approach exploits classic results in the quantum-classical correspondence and reveals interesting connections between the discretization provided by sampling and quantization.

Session S3.2

Invited Talk 11:30
Wed Nov 10

Recent advances in learning quantum states

Srinivasan Arunachalam
IBM

Learning an unknown n-qubit quantum state is a fundamental challenge in quantum computing theory and practice. Information-theoretically, it is well-known that tomography requires exponential in n many copies of an unknown state in order to estimate it upto small trace distance. But a natural question is, are there models of learning where fewer copies suffice and are there interesting classes of states that can be learned with fewer copies? In this talk I will discuss the following results: (1) learning local Hamiltonians on n qubits using poly(n) many samples of the quantum Gibbs state, (2) In the past few years, there have been various learning models introduced to capture the learnability of quantum states; here I will overview many recent results and discuss various equivalences between these learning models. Both these works pave the way towards a more rigorous application of using machine learning techniques to learning quantum states.

Paper 85: Resonant quantum principal component analysis

Zhaokai Li

Principal component analysis (PCA) has been widely adopted to reduce the dimension of data while preserving the information. The quantum version of PCA (qPCA) can be used to analyze an unknown low-rank density matrix by rapidly revealing the principal components of it, i.e., the eigenvectors of the density matrix with the largest eigenvalues. However, because of the substantial resource requirement, its experimental implementation remains challenging. Here, we develop a resonant analysis algorithm with minimal resource for ancillary qubits, in which only one frequency-scanning probe qubit is required to extract the principal components. In the experiment, we demonstrate the distillation of the first principal component of a 4×4 density matrix, with an efficiency of 86.0% and a fidelity of 0.90. This work shows the speedup ability of quantum algorithm in dimension reduction of data and thus could be used as part of quantum artificial intelligence algorithms in the future.

Paper 35: Multi-armed quantum bandits: Exploration vs exploitation when learning properties of quantum states

Josep Lumbreras, Erkka Haapasalo, and Marco Tomamichel

We initiate the study of tradeoffs between exploration and exploitation in online learning of properties of quantum states. Given sequential oracle access to an unknown quantum state, in each round, we are tasked to choose an observable from a set of actions aiming to maximize its expectation value on the state (the reward). Information gained about the unknown state from previous rounds can be used to gradually improve the choice of action, thus reducing the gap between the reward and the maximal reward attainable with the given action set (the regret). We provide various information-theoretic lower bounds on the cumulative regret that an optimal learner must incur, and show that it scales at least as the square root of the number of rounds played. We also investigate the dependence of the cumulative regret on the number of available actions and the dimension of the underlying space. Moreover, we exhibit strategies that are optimal for bandits with a finite number of arms and general mixed states. If we have a promise that the state is pure and the action set is all rank-1 projectors the regret can be interpreted as the infidelity with the target state and we provide some support for the conjecture that measurement strategies with less regret exist for this case.

Paper 33: Provable superior accuracy in machine learned quantum models

Chengran Yang, Andrew Garner, Feiyang Liu, Nora Tischler, Man-Hong Yung, Jayne Thompson, Mile Gu, and Oscar Dahlsten

In building models from big data, there is a constant trade-off between model precision and complexity. Here we show that quantum models enable a more favourable trade-off through machine learning techniques. We develop an algorithm to infer dimensionally constrained quantum models whose accuracy we prove exceeds any classical counterpart. We refer to arxiv: 2105.14434 for more details.

Session S3.3

Training quantum neural networks with an unbounded loss function

Invited Talk 14:30
Wed Nov 10

Maria Kieferova
University of Technology Sydney

Quantum neural networks (QNNs) are a framework for creating quantum algorithms that promises to combine the speedups of quantum computation with the widespread successes of machine learning. A major challenge in QNN development is a concentration of measure phenomenon known as a barren plateau that leads to exponentially small gradients for a range of QNNs models. In this work, we examine the assumptions that give rise to barren plateaus and show that an unbounded loss function can circumvent the existing no-go results. We propose a training algorithm that minimizes the maximal Renyi divergence of order two and present techniques for gradient computation. We compute the closed form of the gradients for Unitary QNNs and Quantum Boltzmann Machines and provide sufficient conditions for the absence of barren plateaus in these models. We demonstrate our approach in two use cases: thermal state learning and Hamiltonian learning. In our numerical experiments, we observed rapid convergence of our training loss function and frequently archived a 99% average fidelity in fewer than 100 epochs.

Paper 83: Parametrized quantum circuits for reinforcement learning

Sofiene Jerbi, Andrea Skolik, Casper Gyurik, Simon Marshall, Hans Briegel, and Vedran Dunjko

We propose two hybrid quantum-classical reinforcement learning models, which we show can be effectively trained to solve several standard benchmarking environments. Moreover, we demonstrate and formally prove the ability of parametrized quantum circuits to solve certain learning tasks that are intractable to classical models, under widely-believed complexity theoretic assumptions.

Paper 48: Structural risk minimization for quantum linear classifiers

Casper Gyurik, Dyon van Vreumingen, and Vedran Dunjko

We provide new insights into how to optimally tune parameterized quantum circuit based machine learning models by balancing the model's complexity against its success at fitting trainingdata (often referred to as structural risk minimization). Specifically, we theoretically quantify how the model's complexity and empirical performance both depend on the rank and Frobenius norm of the observables measured by the model.

Paper 22: Graph neural network initialisation for quantum approximate optimisation

Nishant Jain, *Brian Coyle*, Niraj Kumar and Elham Kashefi

Approximate optimisation is a promising application area for both near-term, and fault-tolerant quantum computers. The quantum approximate optimisation algorithm (QAOA) is one of the foremost algorithms which is suitable to run on noisy intermediate-scale quantum devices, and has seen promising results. In this work, we address two important ingredients in the QAOA for solving the canonical *maxcut* problem. The first is in the initialisation of the circuit parameters, so called 'warm-starting'. Good warm-starting techniques are important to improve the performance and convergence of the algorithm. To this end, we propose graph neural networks (GNNs) for warm-starting the QAOA. This approach has a number of advantages over previous works, in particular, it significantly speeds up initialisation over graph instances, and is capable of generalisation over both problem instance and graph size.

Session S3.4

Understanding machine learning via exactly solvable models

Invited Talk 17:00
Wed Nov 10

Lenka Zdeborova
EPFL

The affinity between statistical physics and machine learning has a long history, I will describe the main lines of this long-lasting friendship in the context of current theoretical challenges and open questions about deep learning. Theoretical physics often proceeds in terms of solvable synthetic models, I will describe the related line of work on solvable models of simple feed-forward neural networks. I will highlight a path forward to capture the subtle interplay between the structure of the data, the architecture of the network, and the learning algorithm. One can then speculate what would be the suitable solvable models for quantum machine learning.

Paper 36: Generation of High Resolution Handwritten Digits with an Ion-Trap Quantum Computer

Manuel Rudolph, Ntwali Toussaint Bashige, Amara Katabarwa, Sonika Johri, Borja Peropadre, and Alejandro Perdomo Ortiz

In this work, we provide the first practical and experimental implementation of a quantum-classical generative algorithm capable of generating high-resolution images of handwritten digits with state-of-the-art gate-based quantum computers.

Paper 63: Learning of Quantum PUFs based on single-qubit gates

Anna Pappa, *Niklas Pirnay* and Jean-Pierre Seifert

Quantum Physical Unclonable Functions (QPUFs) have been proposed as a way to identify and authenticate quantum devices. In this work we examine classical readout QPUFs based on single qubit rotation gates and show that they are not secure - by demonstrating a Machine Learning based attack on a commercial quantum device.

Thursday Nov 11 2021

Session S4.1

The power of data and simple methods for assessing the possibility of quantum advantage in learning

Invited Talk 9:00
Thu Nov 11

Jarrod McClean

Google

The use of quantum computing for machine learning is among the most exciting prospective applications of quantum technologies. However, machine learning tasks where data is provided can be considerably different than commonly studied computational tasks. In this work, we show that some problems that are classically hard to compute can be easily predicted by classical machines learning from data. Using rigorous prediction error bounds as a foundation, we develop a methodology for assessing potential quantum advantage in learning tasks. The bounds are tight asymptotically and empirically predictive for a wide range of learning models. These constructions explain numerical results showing that with the help of data, classical machine learning models can be competitive with quantum models even if they are tailored to quantum problems. We then propose a projected quantum model that provides a simple and rigorous quantum speed-up for a learning problem in the fault-tolerant regime. For near-term implementations, we demonstrate a significant prediction advantage over some classical models on engineered data sets designed to demonstrate a maximal quantum advantage in one of the largest numerical tests for gate-based quantum machine learning to date, up to 30 qubits.

Paper 18: Barren plateaus preclude learning scramblers

Zoe Holmes, Andrew Arrasmith, Bin Yan, Patrick Coles, Andreas Albrecht, and Andrew Sornborger

Scrambling processes, which rapidly spread entanglement through many-body quantum systems, are difficult to investigate using standard techniques, but are relevant to quantum chaos and thermalization. In this Letter, we ask if quantum machine learning (QML) could be used to investigate such processes. We prove a no-go theorem for learning an unknown scrambling process with QML, showing that any variational ansatz is highly probable to have a barren plateau landscape, i.e., cost gradients that vanish exponentially in the system size. We numerically extend our results to approximate scramblers. Our result places generic limits on the learnability of unitaries when lacking prior information.

Paper 14: Long-time simulations with high fidelity on quantum hardware

Joe Gibbs, Kaitlin Gili, Zoe Holmes, Benjamin Commeau, Andrew Arrasmith, Lukasz Cincio, Patrick Coles, and Andrew Sornborger

Moderate-size quantum computers are now publicly accessible over the cloud, opening the exciting possibility of performing dynamical simulations of quantum systems. However, while rapidly improving, these devices have short coherence times, limiting the depth of algorithms that may be successfully implemented. Here we demonstrate that, despite these limitations, it is possible to implement long-time, high fidelity simulations on current hardware. Specifically, we simulate an XY-model spin chain on the Rigetti and IBM quantum computers, maintaining a fidelity of at least 0.9 for over 600 time steps. This is a factor of 150 longer than is possible using the iterated Trotter method. Our simulations are performed using a new algorithm that we call the fixed state Variational Fast Forwarding (fsVFF) algorithm. This algorithm decreases the circuit depth and width required for a quantum simulation by finding an approximate diagonalization of a short time evolution unitary. Crucially, fsVFF only requires finding a diagonalization on the subspace spanned by the initial state, rather than on the total Hilbert space as with previous methods, substantially reducing the required resources. We further demonstrate the viability of fsVFF through large numerical implementations of the algorithm, as well as an analysis of its noise resilience and the scaling of simulation errors.

Session S4.2

Invited Talk 11:00
Thu Nov 11

Generative Models and the Future of Quantum Simulation

Roger Melko
University of Waterloo & Perimeter Institute

One major goal of the current generation of quantum computers is to "simulate" (or emulate) the Hamiltonians found in condensed matter and material systems. Such quantum simulation strategies are particularly important in cases where it is challenging to simulate these systems with traditional computational tools that have been under development for decades. Recently, the rapidly-advancing field of machine learning has introduced a host of new methods suitable for this task, involving neural network architectures and data-driven learning strategies. In this talk, I will discuss the complementary role of experimental and in silico quantum simulations through the lens of generative models, using the example of present-day Rydberg atom quantum computers. In particular, I will illustrate the utility of standard generative models like restricted Boltzmann machines and recurrent neural networks to leverage data from real experiments. I will end by speculating on the future of scientific discovery in quantum many-body simulators that hybridize traditional and data-driven approaches.

Paper 28: Eigenstate extraction with neural-network tomography

Abhijeet Melkani, *Clemens Gneiting*, and Franco Nori

We implement quantum state tomography through the iterative reconstruction of the eigenstates of the experimentally produced mixed states. Our method leverages machine learning through neural-network representations of pure quantum states, which efficiently scale to large quantum systems. We verify our method with experimental data from trapped ion experiments.

Paper 46: Classical variational simulation of the Quantum Approximate Optimization Algorithm

Matija Medvidović and Giuseppe Carleo

We introduce a method to simulate parametrized quantum circuits, an architecture behind many practical algorithms on near-term hardware, focusing on the Quantum Approximate Optimization Algorithm (QAOA). The algorithm can reach unexplored parameter values without requiring large-scale computational resources. Our approach can be used benchmark the next generation of quantum experiments.

Session S4.3

Artificial neural networks for representing quantum many-body states

Invited Talk 15:30
Thu Nov 11

Yusuke Nomura

RIKEN Center for Emergent Matter Science (CEMS)

Quantum many-body systems are the source of various fascinating phenomena such as magnetism and superconductivity. Because the properties of many-body systems are described by quantum many-body wave functions, it is a great challenge to accurately represent quantum many-body wave functions. In this talk, we will show that Boltzmann machines used in machine learning can be useful for this purpose.

Invited Talk 16:15
Thu Nov 11

Advancing classical and quantum variational algorithms for many-body problems

Nobuyuki Yoshioka
University of Tokyo

Both classical and quantum variational algorithms are developing at a surging rate. The key feature underlying two of them is the expressibility of the variational ansatz: neural networks and quantum circuits. In this talk, we first discuss the recent advancement in classical variational algorithms enabled by the use of neural networks, ranging from ground state properties in first-principles calculations, finite-temperature problems, to non-equilibrium properties [1, 2, 3]. After envisioning the coming challenges and future scopes in classical algorithms, we further discuss the major bottleneck for quantum variational algorithms, namely the hardware errors, and introduce an error-mitigation method that is capable of all stochastic, coherent, algorithmic errors [4].

[1] N. Yoshioka, W. Mizukami, and F. Nori, Commun. Phys. 4, 106 (2021).
[2] Y. Nomura*, N. Yoshioka*, and F. Nori, PRL 127, 060601 (2021). (*equal contribution)
[3] N. Yoshioka and R. Hamazaki, PRB 99, 214306 (2019).
[4] N. Yoshioka et al., arXiv:2107.02611.

Session S4.4

Invited Talk 17:00
Thu Nov 11

Rigorous approaches to quantum-assisted machine learning

Jens Eisert
Free University of Berlin

Recent years have enjoyed an increased interest in notions of quantum-assisted machine learning, driven by the hope that quantum algorithms could fare better than classical ones in instances of learning tasks. These advantages could refer to computational speedups, but also to better generalization and other figures of merit. We will have a comprehensive look at what can be rigorously said about aspects of quantum-assisted machine learning. In the first part of the talk, we will discuss the comparative power of classical and quantum learners for generative modelling within the probably approximately correct (PAC) framework, for which we prove a separation between the quantum and classical settings [1]. Following up on this, we will discuss the theoretical properties of model classes derived from parameterized quantum circuits by deriving generalization bounds which depend explicitly on the strategy used for data-encoding. Apart from allowing one to obtain bounds on the performance of trained parameterized quantum circuit models on unseen data, such bounds also facilitate the selection of optimal data-encoding strategies via structural risk minimization [2]. In the light of new findings on the PAC learnability of the output distributions of local quantum circuits, we will discuss how much structure is actually expected to be required to hope for quantum advantages in quantum-assisted machine learning [3]. Technically, we present a negative result - that the

output distributions of super-logarithmic depth Clifford circuits are not sample-efficiently learnable in the statistical query model - as well as a positive result: When directly given access to samples, the output distributions of local Clifford circuits are computationally efficiently PAC learnable by a classical learner. In an outlook, we will briefly discuss other recent work on quantum-assisted machine learning [4-6]. We will conclude that this type of rigorous understanding seems required to aim at fairly assessing the potential of using quantum devices to tackle learning tasks.

[1] On the quantum versus classical learnability of discrete distributions, R. Sweke, J.-P. Seifert, D. Hangleiter, J. Eisert, Quantum 5, 417 (2021).

[2] Encoding-dependent generalization bounds for parameterized quantum circuits, M. C. Caro, E. Gil-Fuster, J. J. Meyer, J. Eisert, R. Sweke, Quantum, in press, arXiv:2106.03880 (2021).

[3] Learnability of the output distributions of local quantum circuits, M. Hinsche, M. Ioannou, A. Nietner, J. Haferkamp, Y. Quek, D. Hangleiter, J.-P. Seifert, J. Eisert, R. Sweke, arXiv:2110.05517 (2021).

[4] Stochastic gradient descent for hybrid quantum-classical optimization, R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Faehrmann, B. Meynard-Piganeau, J. Eisert, Quantum 4, 314 (2020).

[5] Single-component gradient rules for variational quantum algorithms, T. Hubregtsen, F. Wilde, S. Qasim, J. Eisert, arXiv:2106.01388 (2021).

[6] A variational toolbox for quantum multi-parameter estimation, J. J. Meyer, J. Borregaard, J. Eisert, Nature Partner Journal Quantum Information 7 (2021).

Paper 12: Encoding-dependent generalization bounds for parametrized quantum circuits

Matthias C. Caro, *Elies Gil-Fuster*, Jens Eisert, Johannes Jakob Meyer, and Ryan Sweke

We complement recent developments of theoretical performance guarantees for learning models based on parametrized quantum circuits (PQCs). Namely, we prove the first generalization bounds that explicitly take the encoding of the classical input into account. These bounds allow the selection of suitable PQC-based learning models through structural risk minimization.

Paper 13: Generalization in quantum machine learning from few training data

Matthias C. Caro, Hsin-Yuan Huang, Marco Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J. Coles

Modern quantum machine learning (QML) involves optimizing a parameterized quantum circuit on training data, and subsequently makes predictions on testing data. We comprehensively study the generalization performance in QML after training on few data. We support our theory by numerical experiments for quantum phase recognition and unitary compiling.

Paper 1: Simultaneous Perturbation Stochastic Approximation of the Quantum Fisher Information

Julien Gacon, Christa Zoufal, Stefan Woerner, and Giuseppe Carleo

The Quantum Fisher Information matrix (QFIM) is a central metric in promising near-term algorithms, such as Variational Quantum Imaginary Time Evolution. Calculating the QFIM, however, is computationally expensive and scales quadratically with the number of model parameters. Here we propose sampling techniques to approximate the QFIM at a constant cost.

Friday Nov 12 2021

Session S5.1

Computing time-independent molecular properties using variational (quantum) algorithms

Invited Talk 11:00
Fri Nov 12

Wataru Mizukami
Osaka University

Ab initio electronic structure calculations are now firmly established as a vital tool in chemistry. However, there is still a lack of satisfactory approaches for describing systems with many strongly entangled electrons, although the theory for systems with weak electron entanglement has been reaching maturity. In recent years, quantum circuits and neural networks have attracted much attention for their descriptive power to compress the exponential degrees of freedom exhibited by strongly correlated electron systems. The former is called variational quantum eigensolver (VQE). The latter is generally classified as the variational Monte Carlo. The score of these learning methods is energy, but energy is not the only thing of interest in chemistry. It is known that time-independent molecular properties can be calculated as energy derivatives with respect to external parameters. We have been developing methods to compute energy derivatives analytically. The developed methods enable us to calculate the stable structure of molecules, find chemical reaction pathways, estimate anharmonic terms, and locate conical intersections. In this talk, I will introduce the series of results and discuss some issues of the VQE and the direction of future development.

Paper 29: Configurable sublinear circuits for quantum state preparation

Israel F. Araujo, Daniel K. Park, Teresa B. Ludermir, Wilson R. Oliveira, Francesco Petruccione, and *Adenilton J. da Silva*

The theory of quantum algorithms promises unprecedented benefits of harnessing the laws of quantum mechanics for solving certain computational problems. A persistent obstacle to using such algorithms for solving a wide range of real-world problems is the cost of loading classical data to a quantum state. Several quantum circuit-based methods have been proposed for encoding classical data as probability amplitudes of a quantum state. However, they require either quantum circuit depth or width to grow linearly with the data size, even though the other dimension of the quantum circuit grows logarithmically. In this paper, we present a configurable bidirectional procedure that addresses this problem by tailoring the resource trade-off between quantum circuit width and depth. In particular, we show a configuration that encodes an N-dimensional state by a quantum circuit with $O(\sqrt{N})$ width and depth and entangled information in ancillary qubits. We show a proof-of-principle on five quantum computers and compare the results.

Paper 42: Variational Quantum Reinforcement Learning via Evolutionary Optimization

Samuel Yen-Chi Chen, Chih-Min Huang, Chia-Wei Hsing, Hsi-Sheng Goan and Ying-Jer Kao

We present a quantum reinforcement learning framework optimized by evolutionary algorithm. The model combines a quantum-inspired tensor network and variational quantum circuits. We demonstrate via numerical simulation that the proposed model can process input data with dimensions larger than the capacity of existing quantum devices.

Session S5.2

Invited Talk 13:30
Fri Nov 12

Quantum gravity in the lab: matrix quantum mechanics meets quantum computing

Enrico Rinaldi
University of Michigan and RIKEN

Matrix quantum mechanics plays various important roles in theoretical physics, such as a holographic description of quantum black holes. Understanding quantum black holes and the role of entanglement in a holographic setup is of paramount importance for the development of better quantum algorithms (quantum error correction codes) and for the realization of a quantum theory of gravity.

Quantum computing and deep learning offer us potentially useful approaches to study the dynamics of matrix quantum mechanics. For this reason, I will discuss a first benchmark of such techniques to simple models of matrix quantum mechanics.

First, I will introduce a hybrid quantum-classical algorithm in a truncated Hilbert space suitable for finding the ground state of matrix models on NISQ-era devices. Then, I will discuss a deep learning approach to study the wave function of matrix quantum mechanics, even in a supersymmetric case, using a neural network representation of quantum states. Results for the ground state energy will be compared to traditional Lattice Monte Carlo simulations of the Euclidean path integral as a benchmark.

Paper 68: Optimal control of quantum thermal machines with differentiable programming

Ilia Khait, Juan Carrasquilla, and Dvira Segal

Identifying optimal thermodynamic processes has been the essence of thermodynamics since its inception. We harness differentiable programming to optimize finite-time processes in a quantum thermal machine. Overcoming hard physical constraints, our scheme discovers profiles that are superior to previously suggested protocols and finds flaws in a previously employed thermodynamic quantity.

Paper 69: QNLP: Compositional Models of Meaning on a Quantum Computer

Konstantinos Meichanetzidis, Robin Lorenz, Anna Pearson, Alexis Toumi, Giovanni de Felice, Dimitri Kartsaklis, and Bob Coecke

Quantum Natural Language Processing (QNLP) deals with the design and implementation of NLP models intended to be run on quantum hardware. The categorical compositional distributional model of meaning that combines vector space semantics with compositional syntax and grammar, compels a formal analogy of the tensor structure it features with the mathematical structure of quantum theory. This encourages its use for QNLP since grammatical sentences can then naturally be represented as quantum processes. Here we present experimental results for simple sentence classification tasks with small to medium scale datasets, from implementations on noisy quantum computers provided by IBMQ.

Session S5.3

Quantum orthogonal neural networks

Invited Talk 15:05
Fri Nov 12

Iordanis Kerenidis
CNRS/Université Paris Diderot

We will describe some recent results on quantum and classical orthogonal neural networks, namely neural networks where the weight matrices remain orthogonal as a way to avoid redundancy and vanishing gradients. We will show how to implement such orthogonal neural networks with quantum circuits and also how to improve classical back propagation on such networks from previously $O(n^3)$ to the optimal $O(n^2)$. Last, we will describe proof of concept experimental demonstrations of such neural networks for medical image classification.

Paper 27: Regression and Classification with Optimized Random Features: Applications of Exponential Speedup by Quantum Machine Learning without Sparsity and Low-Rankness Assumptions

Hayata Yamasaki, Sathyawageeswar Subramanian, Sho Sonoda, and Masato Koashi

We develop a quantum algorithm for sampling from an optimized probability distribution of random features, in runtime $O(D)$ that is linear in dimension D of input data, so as to significantly reduce and provably minimize the required number of random features for achieving common learning tasks, regression and classification.

Paper 72: Automated NV-centre calibration for quantum internet nodes

Eliska Greplova, Achmed Marif, Matteo Pompili, and Ronald Hanson

Quantum internet is one of the exciting uses of quantum technologies. NV centres have emerged as an excellent platform for distributed entanglement links key for its realisation. In this work we introduce a physics guided reinforcement learning algorithm for their automated tuning and operation completely free of human operator surveillance.