QTML 2020 PROGRAM

Quick Reference Calendar

	Monday 11/09	Tuesday 11/10	Wednesday 11/11	Thursday 11/12
9:00-9:30 EST	**Introduction**	5.1 **Invited:** Nana Liu 5.2 On the Quantum vs Classical Learnability of Discrete Distributions 5.3 A Continuous Variable Born Machine 5.4 Compressed Sensing Tomography for Qudits in Hilbert spaces of non-power-of-two dimensions	**Poster Session 1**	12.1 **Invited:** Maria Kieferova 12.2 Entanglement and Optimization in Hamiltonian Variational Ansatz 12.3 Quantum enhancements for deep reinforcement learning in large spaces
9:30-10:00 EST	1.1 **Invited:** Lei Wang 1.2 Near-Term Quantum-Classical Associative Adversarial Networks 1.3 Quantum versus Classical Generative Modelling in Finance			
10:00-10:30 EST				
10:30-11:00 EST			**Coffee Break**	**Coffee Break**
11:00-11:30 EST	**Coffee Break**	**Coffee Break**	9.1 **Invited:** Maria Schuld 9.2 Machine-learning tools for rapid control, calibration, and characterization of QPUs and other quantum devices 9.3 Visualizing the loss landscape of quantum circuits	13.1 **Invited:** Hendrik Poulsen-Nautrup 13.2 Improved direct state-overlap estimation for NISQ devices 13.3 Fast adiabatic ground state preparation with few measurements
11:30-12:00 EST	2.1 **Invited:** Vedran Dunjko 2.2 Quantum Model Learning Agent 2.3 Variational Fast Forwarding for NISQ Simulations 2.4 Quantum Error Mitigation Via Machine Learning	6.1 **Invited:** Emine Kucukbenli 6.2 Classification and reconstruction of optical quantum states using deep neural networks 6.3 Quantum Computer-Aided design of Quantum Optics Hardware 6.4 Conceptual understanding through efficient inverse-design of quantum optical experiments		
12:00-12:30 EST				
12:30-1:00 EST			**Lunch Break**	**Lunch Break**
1:00-1:30 EST				
1:30-2:00 EST	**Lunch Break**	**Lunch Break**	10.1 **Invited:** Eleanor Rieffel 10.2 Computing Edit Distances with D-Wave 10.3 Quantum-tailored machine-learning characterization of quantum processors	**Industry Session**
2:00-2:30 EST				
2:30-3:00 EST	3.1 **Invited:** Annabelle Bohrdt 3.2 An Interpretable Non-Linear Convolutional Layer for Machine Learning of Multi-Site Correlators 3.3 Quantum computer-aided design: digital quantum simulation of quantum processors 3.4 Neural Network assisted Superconducting Qubit Readout	7.1 **Invited:** Miles Stoudenmire 7.2 Quantum exploration algorithms for multi-armed bandits 7.3 Learning Algorithms for Large-Scale Out-of-Equilibrium Quantum Systems 7.4 Learning quantum many-body states using generative models over the tensor network architecture		
3:00-3:30 EST			11.1 **Invited:** Juan Carrasquilla 11.2 Error mitigation with Clifford quantum-circuit data 11.3 Attention-based Quantum Tomography	
3:30-4:00 EST				14.1 **Invited:** Aram Harrow 14.2 **Invited:** Giacomo Torlai
4:00-4:30 EST				
4:30-5:00 EST	**Coffee Break**	**Coffee Break**	**Coffee Break**	**Coffee Break**
5:00-5:30 EST	4.1 **Invited:** Estelle Inack 4.2 Hardware Efficient Ansätze for Quantum Optimization with Hard Constraints 4.3 Operator Sampling for Shot-frugal Optimization in Variational Algorithms	8.1 **Invited:** Patrick Coles 8.2 Quantum embeddings for machine learning 8.3 Trainability of Quantum Neural Networks	**Poster Session 2**	15.1 **Invited:** Dries Sels 15.2 Data re-uploading for a universal quantum classifier 15.3 Variational Quantum Linear Solver
5:30-6:00 EST				
6:00-6:30 EST				

Daily Program

Monday, November 9, 2020

Time (EST)	Event
9:00 – 9:30am	**Introduction** *Welcome and Opening Remarks*
9:30 – 10:05am	**INVITED** **Lei Wang**, Chinese Academy of Sciences 1.1: *Differentiable Programming Tensor Networks and Quantum Circuits*
10:05 – 10:30am	**Eric Anschuetz,** Massachusetts Institute of Technology 1.2: *Near-Term Quantum-Classical Associative Adversarial Networks*
10:30 – 10:55am	**Brian Coyle,** University of Edinburgh *1.3: Quantum versus Classical Generative Modelling in Finance*
11:00 – 11:30pm	**COFFEE BREAK**
11:30 – 12:05pm	**INVITED** **Vedran Dunjko**, Leiden University 2.1: *Towards quantum advantage via topological data analysis*
12:05 – 12:30pm	**Brian Flynn,** University of Bristol 2.2: *Quantum Model Learning Agent*
12:30 – 12:55pm	**Zoe Holmes,** Los Alamos National Lab 2.3: *Variational Fast Forwarding for NISQ Simulations*
12:55-1:20pm	**Changjun Kim,** KAIST 2.4: *Quantum Error Mitigation Via Machine Learning*
1:30 – 2:30pm	**LUNCH BREAK**
2:30 – 3:05pm	**INVITED** **Annabelle Bohrdt**, Harvard University 3.1: *Quantum Gas Microscopy and Machine Learning*
3:05 – 3:30pm	**Cole Miles,** Cornell University 3.2: *An Interpretable Non-Linear Convolutional Layer for Machine Learning of Multi-Site Correlators*

QTML 2020 Program

3:30 – 3:55pm	**Tim Menke,** Harvard University 3.3: *Quantum computer-aided design: digital quantum simulation of quantum processors*
3:55 – 4:20pm	**Benjamin Lienhard,** Massachusetts Institute of Technology 3.4: *Neural Network assisted Superconducting Qubit Readout*
4:30 – 5:00pm	**COFFEE BREAK**
5:00 – 5:35pm	**INVITED** **Estelle Inack,** Perimeter Institute 4.1: *Variational neural annealing*
5:35 – 6:00pm	**Davide Venturelli,** NASA 4.2: *Hardware Efficient Ansätze for Quantum Optimization with Hard Constraints*
6:00 – 6:25pm	**Andrew Arrasmith,** Los Alamos National Lab 4.3: *Operator Sampling for Shot-frugal Optimization in Variational Algorithms*

Tuesday, November 10, 2020

Time (EST)	**Event**
9:00 – 9:35am	**INVITED** **Nana Liu**, Shanghai Jiao Tong University 5.1: *Security and robustness in quantum machine learning: progress in adversarial quantum learning*
9:35 – 10:00am	**Ryan Sweke,** FU Berlin 5.2: *On the Quantum vs Classical Learnability of Discrete Distributions*
10:00 – 10:25am	**Ieva Čepaitė,** University of Strathclyde 5.3: *A Continuous Variable Born Machine*
10:25 – 10:50am	**Revanth Badveli,** BITS Pilani-Goa Campus *5.4: Compressed Sensing Tomography for qudits in Hilbert spaces of non-power-of-two dimensions*
11:00 – 11:30pm	**COFFEE BREAK**

QTML 2020 Program

11:30 – 12:05pm	**INVITED** **Emine Kucukbenli**, Harvard University 6.1: *Bypassing Quantum Mechanics with Neural Networks for Atomistic Modelling*
12:05 – 12:30pm	**Shahnawaz Ahmed,** Chalmers University of Technology 6.2: *Classification and reconstruction of optical quantum states using deep neural networks*
12:30 – 12:55pm	**Jakob Kottmann,** University of Toronto 6.3: *Quantum Computer-Aided design of Quantum Optics Hardware*
12:55-1:20pm	**Mario Krenn,** University of Toronto 6.4: *Conceptual understanding through efficient inverse-design of quantum optical experiments*
1:30 – 2:30pm	**LUNCH BREAK**
2:30 – 3:05pm	**INVITED** **Miles Stoudenmire**, Flatiron Institute 7.1: *Perspective on Tensor Networks for Machine Learning*
3:05 – 3:30pm	**Daochen Wang,** University of Maryland 7.2: *Quantum exploration algorithms for multi-armed bandits*
3:30 – 3:55pm	**Eliska Greplova,** ETH Zurich 7.3: *Learning Algorithms for Large-Scale Out-of-Equilibrium Quantum Systems*
3:55 – 4:20pm	**Khadijeh Najafi,** Harvard-Caltech 7.4: *Learning quantum many-body states using generative models over the tensor network architecture*
4:30 – 5:00pm	**COFFEE BREAK**
5:00 – 5:35pm	**INVITED** **Patrick Coles,** Los Alamos National Lab 8.1: *Promises and Challenges of Variational Quantum Algorithms*
5:35 – 6:00pm	**Aroosa Ijaz,** Xanadu, University of Waterloo 8.2: *Quantum embeddings for machine learning*
6:00 – 6:25pm	**Marco Cerezo,** Los Alamos National Laboratory 8.3: *Trainability of Quantum Neural Networks*

QTML 2020 Program

Wednesday, November 11, 2020

Time (EST)	Event
9:00 – 10:30am	**POSTER SESSION 1**
10:30 – 11:00am	**COFFEE BREAK**
11:00 – 11:35am	**INVITED** **Maria Schuld**, University of KwaZulu-Natal, Xanadu 9.1: *Which functions can quantum machine learning models learn?*
11:35 – 12:00pm	**Shai Machnes,** Saarland University 9.2: *Machine-learning tools for rapid control, calibration, and characterization of QPUs and other quantum devices*
12:00 – 12:25pm	**Kathleen Hamilton,** Oak Ridge National Laboratory 9.3: *Visualizing the loss landscape of quantum circuits*
12:30 – 1:30pm	**LUNCH BREAK**
1:30 – 2:05pm	**INVITED** **Eleanor Rieffel**, NASA Ames 10.1: *Utilizing NISQ devices for evaluating quantum algorithms*
2:05 – 2:30pm	**Riccardo Mengoni,** University of Verona 10.2: *Computing Edit Distances with D-Wave*
2:30 – 2:55pm	**Élie Genois,** Institut quantique, Université de Sherbrooke 10.3: *Quantum-tailored machine-learning characterization of quantum processors*
3:00 – 3:35pm	**INVITED** **Juan Carrasquilla**, Vector Institute 11.1: *Simulating quantum systems with probabilistic models*
3:35 – 4:00pm	**Piotr Czarnik,** Los Alamos National Laboratory 11.2: *Error mitigation with Clifford quantum-circuit data*
4:00 – 4:25pm	**Peter Cha,** Cornell University 11.3: *Attention-based Quantum Tomography*
4:30 – 5:00pm	**COFFEE BREAK**
5:00 – 6:30pm	**POSTER SESSION 2**

QTML 2020 Program

Thursday, November 12, 2020

Time (EST)	Event
9:00 – 9:35am	**INVITED** **Maria Kieferova**, University of Technology Sydney 12.1: *Entanglement Induced Barren Plateaus*
9:35 – 10:00pm	**Cunlu Zhou,** University of Toronto 12.2: *Entanglement and Optimization in Hamiltonian Variational Ansatz*
10:00 – 10:25pm	**Sofiene Jerbi,** University of Innsbruck 12.3: *Quantum enhancements for deep reinforcement learning in large spaces*
10:30 – 11:00am	**COFFEE BREAK**
11:00 – 11:35am	**INVITED** **Hendrik Poulsen-Nautrup**, University of Innsbruck 13.1: *Operationally meaningful representations of physical systems in neural networks*
11:35 – 12:00pm	**Leonardo Guerini,** Federal University of Rio de Janeiro 13.2: *Improved direct state-overlap estimation for NISQ devices*
12:00 – 12:25pm	**Benjamin F. Schiffer,** Max Planck Institute of Quantum Optics 13.3: *Fast adiabatic ground state preparation with few measurements*
12:30 – 1:30pm	**LUNCH BREAK**
1:30 – 3:30pm	**INDUSTRY SESSION** PANELISTS: **Christa Zoufal,** IBM Quantum **Martin Leib**, Volkswagen **Murphy Yuezhen Niu**, Google **Nathan Killoran**, Xanadu **Philip Makotyn**, Honeywell **Stefan Leichenauer**, X the Moonshot Factory
3:30 – 4:05pm	**INVITED** **Aram Harrow**, Massachusetts Institute of Technology 14.1: *Small quantum computers and large classical data sets*

QTML 2020 Program

4:05 – 4:40pm	**INVITED** **Giacomo Torlai**, AWS Center for Quantum Computing 14.2: *Learning quantum channels with tensor networks*
4:40 – 5:00pm	**COFFEE BREAK**
5:00 – 5:35pm	**INVITED** **Dries Sels**, Harvard University 15.1: *Quantum Sampling algorithm for Near-Term Devices*
5:35 – 6:00pm	**Alba Cervera-Lierta,** University of Toronto 15.2: *Data re-uploading for a universal quantum classifier*
6:00 – 6:25pm	**Lukasz Cincio,** University of Barcelona 15.3: *Variational Quantum Linear Solver*

Abstract Library

1.1	**Differentiable Programming Tensor Networks and Quantum Circuits** Lei Wang

Differentiable programming is a fresh programming paradigm that emerges from deep learning but is not only limited to training neural networks. In differentiable programming, one composes differentiable components into an adaptive program and learns the program with gradient based optimization. I will show scientific applications of differentiable programming on statistical and quantum physics problems. Finally, I will discuss differentiable programming on quantum circuits which is relevant to applications of near-term quantum devices.

1.2	**Near-Term Quantum-Classical Associative Adversarial Networks** Eric Anschuetz and Cristian Zanoci

We introduce a new hybrid quantum-classical adversarial machine learning architecture called a quantum-classical associative adversarial network (QAAN). This architecture consists of a classical generative adversarial network with a small auxiliary quantum Boltzmann machine that is simultaneously trained on an intermediate layer of the discriminator of the generative network. We numerically study the performance of QAANs compared to their classical counterparts on the MNIST and CIFAR-10 data sets, and show that QAANs attain a higher quality of learning when evaluated using the Inception score and the Fréchet Inception distance. As the QAAN architecture only relies on sampling simple local observables of a small quantum Boltzmann machine, this model is particularly amenable for implementation on the current and next generations of quantum devices.

1.3	**Quantum versus Classical Generative Modelling in Finance** Brian Coyle, Maxwell Henderson, Justin Chan Jin Le, Niraj Kumar, Marco Paini and Elham Kashefi

Finding a concrete use case for quantum computers in the near term is still an open question, with machine learning typically touted as one of the first fields which will be impacted by quantum technologies. In this work, we investigate and compare the capabilities of quantum versus classical models for the task of generative modelling in machine learning. We use a real world financial dataset consisting of correlated currency pairs and compare two models in their ability to learn the resulting distribution - a restricted Boltzmann machine, and a quantum circuit Born machine. We provide extensive numerical results indicating that the simulated Born machine always at least matches the performance of the Boltzmann machine in this task, and demonstrates superior performance as the model scales. We perform experiments on both simulated and physical quantum chips using the Rigetti forest platform, and also are able to partially train the largest instance to date of a quantum circuit Born machine on quantum hardware. Finally, by studying the entanglement capacity of the training Born machines, we find that entanglement typically plays a role in the problem instances which demonstrate an advantage over the Boltzmann machine.

2.1	**Towards quantum advantage via topological data analysis** Vedran Dunjko
A significant aspect of quantum machine learning (QML) relying on quantum linear algebra ideas has been hailed as a potential source of exponential speedups over classical algorithms. However, a recent series of “dequantization” results, which use clever classical sampling methods, have removed the possibility of strict exponential speed-ups for many machine learning and related quantum-linear-algebraic approaches. However, not all such algorithms can directly be de-quantized using these recent classical methods. One example of such a QML algorithm, that has not (yet) been dequantized is the quantum algorithm for topological data analysis, introduced by Lloyd, Garnerone and Zanardi. Can exponential separations for this problem be proven, at least up to some more established complexity-theoretic assumptions? Further, does this algorithm have a chance of being extended to other problems and run on near-term devices? In this talk we will provide evidence that the problem of topological data analysis, and certain other related more general problems, are truly out of reach for classical computers. Towards the end of the talk we will also consider the suitability of the quantum topological data analysis algorithm for near-term implementations, and discuss the possibility of it being one of the first quantum "NISQ killer apps".	
2.2	**Quantum Model Learning Agent** Brian Flynn, Andreas Gentile, Sebastian Knauer, Stefano Paesani, Nathan Wiebe, Christopher Granade, John Rarity, Raffaele Santagati and Anthony Laing
Hamiltonian operators provide a compact description of physical systems, modelling essential properties including the interactions within the system and its energy, and predicting its behaviour. Knowing the Hamiltonian of quantum systems is then essential for studying their physics and for developing quantum technologies. However, inferring Hamiltonians of quantum systems from observations is a difficult problem, because noise and unfavourable scaling limit our ability to formulate and test models reliably. Here, we demonstrate an unsupervised machine learning algorithm which retrieves Hamiltonian models from quantum systems, represented in an interpretable form. By learning human-readable models with high predictive power, our method offers a key tool in automating the investigation of uncharacterised quantum systems, which is a crucial challenge for developing quantum technologies.	
2.3	**Variational Fast Forwarding for NISQ Simulations** Cristina Cirstoiu, Zoe Holmes, Benjamin Commeau, Joseph Iosue, Lukasz Cincio, Patrick Coles and Andrew Sornborger
Quantum simulation has the potential to transform fields from materials science to quantum chemistry. This could lead to improved pharmaceutical development, new catalysts to capture carbon in the atmosphere, and higher temperature superconductors for more effective power transmission. However, while quantum hardware is rapidly reaching the stage where it can outperform classical supercomputers, we remain in the noisy intermediate-scale quantum (NISQ) era in	

which devices are prone to errors and susceptible to decoherence. The resulting short coherence times of near term quantum computers restrict the length of times that may be successfully simulated thus limiting the usefulness of quantum simulations. Here we present a hybrid quantum-classical algorithm, called Variational Fast Forwarding (VFF), to address the challenge of how to perform genuinely useful simulations with NISQ devices. The number of gates required for a VFF simulation does not scale with the length of time to be simulated and thus enables simulations of long times using a fixed number of gates. To achieve this, VFF approximately diagonalises a Trotter approximation of the short-time evolution using a hybrid quantum-classical variational method, and in turn this diagonalisation allows for arbitrary time simulations with the same circuit structure. We test the capability of VFF through extensive numerical simulations of the Hubbard, Ising, and Heisenberg models, and we implement VFF on Riggeti's quantum computer to demonstrate simulation beyond the coherence time. In contrast to prior work on variational quantum simulations, we further provide a rigorous analysis of VFF's simulation errors. We find that the simulation error of VFF scales at worst linearly in the fast-forwarded simulation time. This enables us to derive a natural termination condition for the algorithm which is framed in terms of the required average fidelity of the simulation. We further present two refinements of the VFF algorithm. The first, called Hamiltonian Fast Forwarding (HFF), reduces the error induced by the initial Trotterisation by directly diagonalising the system's Hamiltonian, thus improving the accuracy of the simulation. While VFF and HFF can be applied to simulate any quantum state, in many cases one is interested in simulating a particular initial quantum state. Our second refinement, called fixed-state VFF (fsVFF), is tailored to this situation. The cost function for fsVFF requires fewer qubits making it yet more NISQ friendly.

2.4	**Quantum Error Mitigation Via Machine Learning** Changjun Kim, Daniel K. Park and June-Koo Kevin Rhee

"The main obstacle to the development of full-fledged quantum computers that promise revolutionary opportunities is computational errors that occur when theoretical ideas are implemented on real quantum devices. The theory of quantum error correction (QEC) and fault-tolerance guarantees scalable quantum computation, but the computational resource overhead is non-negligible. Moreover, the typical error rate of current quantum devices are near or above the fault-tolerance threshold value, hindering the application of quantum error correction. With this background, the idea of quantum error mitigation (QEM) emerged recently. Unlike QEC, QEM does not aim to fully remove the entropy increased by unwanted interaction with the environment to recover the logical state. Instead, it aims to merely improve the accuracy of estimating the final answer in a given computational task without having to encode logical quantum state in a multi-qubit entangled state. Since QEM does not require extra quantum resource, it is expected to improve the quantum computation to some extent even when the error rate is above the fault-tolerance threshold value. Thus, QEM is an excellent fit for improving the performance of the noisy

intermediate-scale quantum (NISQ) computing. In this work, we propose a QEM technique based on two machine learning methods. The machine learning aims to learn how to adjust the probabilities estimated by the projective measurement in the computational basis. The training is done by repeating the evolution of a two-qubit system under a known quantum circuit randomly generated with 1 to 10 single- and two-qubit gates and comparing the actual outcome to the expected outcome. After the training, when a new arbitrary quantum circuit is executed, the QEM adjusts the measurement statistics accordingly. The QEM is tested on an arbitrary two-qubit quantum circuit with 1 to 20 gates, both with simulation using realistic error models and with the IBM quantum cloud platform. We compared the error reduction achieved by the machine learning based QEM with that by an analytical method. Deep neural network (DNN) and Concatenated DNN are used for machine learning. We propose an effective design of DNN that can learn about long sequences of gates without requiring full knowledge of the permutations of the gates. In DNN all layers are fully connected and ReLu is used as activation function. There are two hidden layers and measurement outcome and the number of gates are used as input. For concatenated DNN, batch normalization is added for the number of gates to normalize the input in order to improve the performance of DNN. Concatenated DNN shows the best result among all mitigation methods when simulated with a realistic error model. It reduces the error by at least a factor of 0.65 for simulated data and by at least a factor of 0.46 for data from the five-qubit IBM quantum device. An intriguing finding is that the proposed method performs better with higher error rate, even higher than the fault-tolerant threshold."

3.1	**Quantum Gas Microscopy and Machine Learning** Annabelle Bohrdt

Quantum gas microscopes for ultracold atoms provide high-resolution real-space snapshots of the quantum state of complex many-body systems. Each measurement collapses the quantum state, such that only a projection of it can be accessed. Repeated projective measurements yield a plethora of data, which in the past have mostly been analyzed to obtain conventional observables such as one- and two-point correlation functions, which are also traditionally measured in solids. However, measurements performed in quantum gas microscopes contain considerably more information. Therefore, the need arises for new methods to analyze the data that take all available information into consideration and hence use the capabilities of quantum gas microscopes to their full extent. In this talk, I will present our results on using machine learning techniques to evaluate and characterize snapshots from quantum gas microscopy experiments realizing two paradigmatic models: the Bose- and Fermi-Hubbard model. In particular, we study dynamics in the one-dimensional Bose-Hubbard model with a disorder potential, which realizes a many-body localized phase at sufficiently strong disorder. Moreover, we compare snapshots of the two-dimensional Fermi–Hubbard model to two theoretical approaches: a doped quantum spin liquid state of resonating valence bond type, and the geometric string theory, describing a state with hidden spin order. Using machine learning techniques to

analyze snapshots of the quantum many-body state takes into account all available information without a potential bias by the choice of an observable.

3.2	**An Interpretable Non-Linear Convolutional Layer for Machine Learning of Multi-Site Correlators** Cole Miles, Annabelle Bohrdt, Ruihan Wu, Kilian Q. Weinberger, Eugene Demler and Eun-Ah Kim

Machine learning approaches have recently found application in processing Quantum Gas Microscopy experimental data, where they have been found to be able to perform beyond the level that measurements of simple physical observables would allow. However, whatever physics these models are learning has thus far been locked behind the ""black box"" of their complexity. Here, we present a set of novel nonlinear operations used to construct a new convolutional layer that discovers features which are directly interpretable in terms of physical observables. This layer is applicable to general lattice data, and lends itself well to construction of simple, end-to-end interpretable architectures while maintaining performance relative to more complex models. In anticipation of the next generation of quantum gas microscopy experiments, we demonstrate this new architecture on sets of simulated snapshots sampled from two candidate theories approximating the doped Fermi-Hubbard model. From the trained networks, we discover that the network learns meaningful high-order spin-hole correlators which significantly differ between the theories. This architecture provides a more rigorous way of comparing future experimental data to theoretical predictions, and may be generalized to study other systems.

3.3	**Quantum computer-aided design: digital quantum simulation of quantum processors** Tim Menke, Thi Ha Kyaw, Sukin Sim, Nicolas Sawaya, William Oliver, Gian Giacomo Guerreschi and Alán Aspuru-Guzik

"With the increasing size of quantum processors, sub-modules that constitute the processor hardware will become too large to accurately simulate on a classical computer. Therefore, one would soon have to fabricate and test each new design primitive and parameter choice in time-consuming coordination between design, fabrication, and experimental validation. Here we show how one can design and test the performance of next-generation quantum hardware – by using existing quantum computers. Focusing on superconducting transmon processors as a prominent hardware platform, we compute the properties of individual and coupled transmons. We develop explicit quantum algorithms for hardware simulation, provide resource counts, and give a quantitative estimate for when quantum hardware simulation capabilities will surpass those of classical simulations. On a quantum computer, the transmon Hamiltonian can be encoded in the Hilbert space of four qubits after truncation to the relevant basis states. The qubit count then grows linearly in the transmon processor size, as opposed to exponentially if one were to use classical computing resources. We demonstrate how previously established encoding

techniques [1] are used to convert the multi-level superconducting circuit Hamiltonian into Pauli operators. The energy spectra of the transmon circuits are computed by variational hybrid quantum-classical algorithms that are well-suited for near-term noisy quantum computers. In addition to employing the variational quantum eigensolver (VQE) to find the ground state, the recently proposed variational quantum deflation (VQD) algorithm [2] is used to estimate excites states. A highly expressible and hardware-efficient variational ansatz was developed in order to reliably determine the first few eigenstates of single- and two-transmon systems, which are encoded in four and eight qubits, respectively. The resulting spectra are of practical relevance, as they can pinpoint avoided energy level crossings for gate operations as well as shed light on unwanted cross-talk arising from spurious capacitances in the circuit. The developed variational ansatz can be extended to the simulation of multiple transmons in a processor sub-module. We find that the depth of the variational circuit scales linearly and the two-qubit gate count scales quadratically with the number of transmons. The analysis is based on the capabilities of ion trap quantum computers and shows that the algorithm can be readily applied on near-term quantum computers. Therefore, our methods pave a new way towards designing candidate quantum processors when the demands of calculating sub-module properties exceed the capabilities of classical computing resources.

3.4	**Neural Network assisted Superconducting Qubit Readout** Benjamin Lienhard, Antti Vepsäläinen, Luke C. G. Govia, Cole R. Hoffer, Jack Y. Qiu, Diego Ristè, Matthew Ware, David Kim, Roni Winik, Alexander J. Melville, Bethany M. Niedzielski, Jonilyn L. Yoder, Guilhem Ribeill, Thomas A. Ohki, Hari Krovi, Terry P. Orlando, Simon Gustavsson and William D. Oliver

Qubit-state readout is a significant error source in contemporary quantum processors. As a consequence of the increasing number of qubits constituting today's devices, readout discrimination has become a computationally intensive, multi-state classification problem. For a single superconducting qubit connected via a resonator to a readout bus, a linear matched filter is sufficient for high-fidelity readout. However, it is more resource efficient to frequency-multiplex multiple qubits on a single readout line and, in this case, the use of multiple, single-qubit matched filters becomes increasingly challenging to implement with high fidelity. Here, we present an alternative approach to multi-qubit readout discrimination based on neural networks. We consider different types of neural network architectures and their relative performance compared to current readout methods when applied to experimental single and multi-qubit readout data. We demonstrate an improvement in the multi-qubit assignment fidelity of up to 4% per qubit using neural networks, compared to conventional matched filters, primarily due to the compensation of system dependent non-idealities such as readout crosstalk. The general approach demonstrated here is applicable to other qubit modalities that leverage frequency-multipexed multi-qubit readout.

4.1	**Variational neural annealing** Estelle Inack

Many combinatorial optimization problems relevant to computer science, computational biology and physics can be tackled with simulated annealing, which is a powerful framework for optimizing the properties of complex systems through the lens of statistical mechanics. However, simulated annealing and its quantum counterpart, simulated quantum annealing, are traditionally implemented via Markov chain Monte Carlo, often displaying slow convergence to optimal solutions for challenging optimization problems. Here we combine the variational principle in classical and quantum physics with recurrent neural networks, whose dynamics are naturally devoid of slow Markov chains to accurately emulate annealing in its classical and quantum formulations. We find that a variational implementation of classical annealing is not only superior to its quantum analog in terms of speed of convergence and accuracy of solutions but also outperforms traditional simulated annealing and simulated quantum annealing on several prototypical spin glass models as the annealing time is increased. In addition, through the use of the sign-problem-free variational Monte Carlo method, our framework can emulate quantum annealing with non-stoquastic drivers at moderately large system sizes, thus providing a useful tool to benchmark the next generation of quantum annealing devices which implement non-stoquastic Hamiltonians.

4.2	**Hardware Efficient Ansätze for Quantum Optimization with Hard Constraints** Ryan LaRose, Eleanor Rieffel and Davide Venturelli

We introduce multiple parametrized circuit ansätze and present the results of a numerical study comparing their performance with a standard Quantum Alternating Operator Ansatz (QAOA) (Algorithms 12 (2), 34, 2019) approach. The ansätze were inspired by compilation considerations, with the aim of keeping some of the mixing and phase separation intuitions of QAOA while reducing the depth of the circuits, which would have a significant advantage when run on NISQ hardware. They are particularly suitable for quantum processors, such as several superconducting processors, that natively implement an XY interaction Hamiltonian between neighboring qubits. The methods are tested on random instances of a weighted quadratic binary constrained optimization problem – fully connected – that feature a search space over feasible solutions of constant Hamming weight, so grows as 2^N/Sqrt(N) with the number of qubits N. The results indicate that Ansätze that contains information on the cost function Hamiltonian as part of their initial structure perform better than their counterparts that do not include cost function weights. Moreover, for all the parameter setting strategies used, and for two performance metrics explored, the average performance achieved by the QAOA method is best but effectively matched by the one obtained by a “mixer-phaser” ansatz (QAMPA) that can be compiled in less than half-depth of standard XY-QAOA under reasonable assumptions. This mixer-phaser ansatz can be tested experimentally at the same scale as the simulations we used to obtain our results on a variety of current and near term quantum processors and can be generalized to more complex constrained

optimization problems, such as problems common in the scheduling, logistics, finance, telecommunication domains.

4.3	**Operator Sampling for Shot-frugal Optimization in Variational Algorithms** Andrew Arrasmith, Lukasz Cincio, Rolando Somma and Patrick Coles

Quantum chemistry is a near-term application for quantum computers. This application may be facilitated by variational quantum-classical algorithms (VQCAs), although a concern for VQCAs is the large number of measurements needed for convergence, especially for chemical accuracy. Here we introduce a strategy for reducing the number of measurements (i.e., shots) by randomly sampling operators from the overall Hamiltonian. In particular, we employ weighted sampling, which is important when the weights on the operators are highly non-uniform, as is typical in chemistry. We integrate this strategy with an adaptive optimizer developed recently by our group to construct an improved optimizer called Rosalin (Random Operator Sampling for Adaptive Learning with Individual Number of shots). Rosalin implements stochastic gradient descent while adapting the shot noise for each partial derivative and randomly assigning the shots amongst the h_i according to a weighted distribution. We implement this and other optimizers to find the ground states of molecules H2, LiH, and BeH2, without and with quantum hardware noise, and Rosalin outperforms other optimizers in most cases.

5.1	**Security and robustness in quantum machine learning: progress in adversarial quantum learning** Nana Liu

The central questions in adversarial quantum learning lies in identifying the security vulnerabilities of quantum machine learning, how to protect learning algorithms against security attacks and how to leverage quantum features for security and other robustness advantages. This is relevant not only for real-life deployment of quantum machine learning algorithms in future quantum networks but is also an important platform to ask fundamental questions in machine learning with quantum data. In this talk, I'll introduce the emerging area of adversarial quantum learning and outline its recent progress. I'll present some key open questions and connections between this area and more traditional topics in quantum information theory, then suggest methods forward to seek alternative measures of quantum advantage for machine learning in terms of resistance to adversaries.

5.2	**On the Quantum vs Classical Learnability of Discrete Distributions** Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter and Jens Eisert

Since its introduction, Valiant's model of Probably Approximate Correct (PAC) learning has provided a fruitful framework for studying both the computational and statistical aspects of machine learning, and for comparing in a rigorous way the comparative power of quantum and classical learning algorithms. However, the majority of previous work in this framework has been focused on the supervised learning setting, and as of yet there are no rigorous results addressing the potential advantages of quantum learners for unsupervised learning. In order to address this question we

present a rigorous study of the comparative power of classical and quantum learners for generative modelling, within the PAC framework. More specifically we consider the following task: Given samples from some unknown discrete probability distribution, output with high probability an efficient algorithm for generating new samples from a good approximation to the original distribution. Our primary result is the explicit construction of a class of discrete probability distributions which, under the decisional Diffie-Hellman assumption, is provably not efficiently PAC learnable by a classical generative modelling algorithm, but for which we construct an efficient quantum learner. This class of distributions therefore provides a concrete example of a generative modelling problem for which quantum learners exhibit a provable advantage over classical learning algorithms. In order to present this result we develop the PAC framework for generative modelling in a systematic way, which also allows us to describe and discuss a variety of natural open questions concerning the power of near-term generative modelling algorithms, the power of learning from quantum samples and the classical hardness of learning distributions which cannot be classically sampled from.

5.3	**A Continuous Variable Born Machine** Ieva Čepaitė, Brian Coyle and Elham Kashefi

Generative Modelling has become a promising use case for near term quantum computers. In particular, due to the fundamentally probabilistic nature of quantum mechanics, quantum computers naturally model and learn probability distributions, perhaps more efficiently than can be achieved classically. The Born machine is an example of such a model, easily implemented on near term quantum computers. However, in its original form, the Born machine only naturally represents discrete distributions. Since probability distributions of a continuous nature are commonplace in the world, it is essential to have a model which can efficiently represent them. Some proposals have been made in the literature to supplement the discrete Born machine with extra features to more easily learn continuous distributions, however, all invariably increase the resources required to some extent. In this work, we present the continuous variable Born machine, built on the alternative architecture of continuous variable quantum computing, which is much more suitable for modelling such distributions in a resource-minimal way. We provide numerical results indicating the models ability to learn both quantum and classical continuous distributions, including in the presence of noise.

5.4	**Compressed Sensing Tomography for Qudits in Hilbert spaces of non-power-of-two dimensions** Revanth Badveli, Vinayak Jagadish, Srikanth Radhakrishna and Francesco Petruccione

The matrix generalizations of Compressed Sensing (CS) were adapted to Quantum State Tomography (QST) previously by Gross et al. [Phys. Rev. Lett. 105, 150401 (2010)], where they consider the tomography of n spin-1/2 systems. For the density matrix of dimension d = 2^n and rank r with r << 2^n, it was shown that randomly chosen Pauli measurements of the order O[dr \log(d)^2] are enough to fully

reconstruct the density matrix by running a specific convex optimization algorithm. However, these results utilized the low operator-norm of the Pauli operator basis, which are available only in power-of-two dimensional Hilbert spaces. In the present work [Phys. Rev. A 101, 062328 (2020)], we propose an alternate CS-QST protocol for states in Hilbert spaces of non-power-of-two dimensions, which still achieves the bounds on number of measurement settings O[dr \log(d)^2] presented in [Phys. Rev. Lett. 105, 150401 (2010)]. In this alternate protocol, we use a unitary operator W to "move" the quantum information from a d dimensional system to a d_1 dimensional ancilla, where d_1 is a power of two. We prove that choosing the optimal value for d_1 and performing the standard CS-QST protocol using simple Pauli measurements on the ancilla will guarantee full recovery from O[dr \log(d)^2] measurements. We show that the unitary operator W can be efficiently implemented using only poly[log(d)^2]$ single qubit gates at most, which is relatively a small overheard compared to the cost of CS-QST protocol. For states in Hilbert spaces of non-power-of-two dimensions, one may consider performing the standard CS-QST protocol using the SU(d) operators. We point out that the SU(d) operators, owing to their high operator norm, do not provide a significant savings in the number of measurement settings required for successful recovery of all rank-r states. We use numerical simulations to show that the proposed alternate approach outperforms the one using SU(d) operators.

6.1	**Bypassing Quantum Mechanics: Atomistic Modelling with Neural Networks** Emine Kucukbenli

In the 90s, the atomistic modelling of materials witnessed a paradigm shift thanks to the increased computational capability of the era, joining forces with an elegant quantum mechanical theory of many-electron systems called Density Functional Theory. Today we may be witnessing the early days of another paradigm shift: with increased computational resources -this time in the form of software libraries- coupled with not-so-elegant but practical adaptation of neural network approaches, we are attempting to bypass the quantum mechanical model to enable affordable simulations of materials. Using examples from 2D, few-layer, and 3D carbon allotropes, I will give a brief overview of how neural networks are utilized for this purpose, where they have led promising outcomes, raised new questions or revealed remaining challenges.

6.2	**Classification and reconstruction of optical quantum states using deep neural networks** Shahnawaz Ahmed, Carlos Sánchez Muñoz, Franco Nori and Anton Frisk Kockum

"Machine-learning techniques for quantum state tomography can give significantly faster results, e.g., by using adaptive measurements to avoid redundant data acquisition, or by finding efficient parameterisations of a quantum state to escape the ""curse of dimensionality"". Here, we explore a similar idea for optical quantum states by using deep neural networks (DNNs). We train a DNN to classify different optical quantum states, e.g., cat or thermal states, with high accuracy. We test the

robustness of our DNN classification against simulated noisy data and show that it is possible to train a DNN to recognise different classes of states, including bosonic codes such as binomial states, when they exist as mixed states. Further, our DNN approach can reconstruct the full density matrix of the quantum state from noisy measurement data. We show two possible methods for reconstruction - one involving training on simulated data and the other reconstruction from scratch using a single instance of previously unseen data. The training of the DNN is performed by standard gradient-based optimization routines such as Adam in both cases. We show results on experimental data and compare our reconstruction method to iterative maximum likelihood estimation. Our work thus explores the use of DNNs for fast and efficient real-time classification and reconstruction of optical quantum states in the lab."

6.3	**Quantum Computer-Aided design of Quantum Optics Hardware** Jakob Kottmann, Mario Krenn, Thi Ha Kyaw, Sumner Alperin-Lea and Alán Aspuru-Guzik

Photonic devices are one of the key quantum technologies. Due to the non-interacting nature of photons, the design of photonic state preparation setups are challenging even for conceptually simple quantum states. Computational tools are playing an increasingly larger role in photonic experiment design. Currently, classical algorithms used for this purpose are fundamentally limited by the size of the underlying photonic setup due to its quantum mechanical nature. We show explicitly how to map quantum optical setups to gate based quantum circuits which grants the possibility to optimize and verify their actions with the help of digital quantum computers. In combination with classical topological optimizers, this provides a path towards the computational design of large photonic devices in the future. Our algorithms are publicly available and implemented with the open-source package tequila that provides an abstract high-level framework for the development of novel quantum algorithms. We illustrate how such a framework is able to mutually improve algorithmic development within different fields like quantum chemistry and quantum optics.

6.4	**Conceptual understanding through efficient inverse-design of quantum optical experiments** Mario Krenn, Jakob S. Kottmann, Nora Tischler and Alan Aspuru-Guzik

"The design of quantum experiments can be challenging for humans. In experimental quantum optics, A.I. methods have therefore been introduced to solve the inverse-design problem. While some computer-designed experiments have been successfully demonstrated in laboratories, these algorithms generally are slow, require a large amount of data or work for specific platforms that are difficult to generalize. Here we present *theseus*, an efficient algorithm for the design of quantum experiments, which we use to solve several open questions in experimental quantum optics. The algorithm's core is a physics-inspired, graph-theoretical representation of quantum states, which makes it significantly faster than previous comparable approaches. The improvement in speed shows that *theseus* is ready to go beyond benchmarks, and be applied to the *discovery of new scientific targets* and -- more importantly -- to the

development of new scientific insights and understanding. *Scientific understanding* is essential to the epistemic aims of science , but rarely addressed in applications of artificial intelligence to the natural sciences. In the philosophy of science, pragmatic criteria have been found for *scientific understanding*, in particular by de Regt's award-winning work. He describes that scientists can understand a phenomenon *if they can recognise qualitatively characteristic consequences without performing exact calculations*. We connect this criterion to our discoveries: We discover the first implementation of high-dimensional GHZ states, important classes of heralded photonic states and new experimental configurations of previously unknown transformations. In all of these cases, we identify the underlying conceptual cores of the solution. We use them to generalize the solutions to many other, large classes of states and transformations -- without additional calculations. Connecting to de Regt's criterion that argue that scientific understanding is connected with the skill to use concepts fruitfully, *without exact calculations*, our algorithm has been the source of scientific understanding for multiple instances. This is in stark contrast to typical A.I. applications in the natural sciences, where the solution of a parameter optimisation is the final product, without the intention of discovering new scientific ideas. Hence, in a broader sense, the ability of our algorithm goes beyond simple optimisation, and enters the realm of providing scientific insights and allowing for scientific understanding. Finally, we believe that insights from the philosophy of science can be productive guiding principles in the development of new A.I. algorithms for physics."

7.1	**Perspective on Tensor Networks for Machine Learning** Miles Stoudenmire

The use of quantum-inspired tensor networks for machine learning is a growing subfield, with frontiers not only in applied machine learning, but also topics like classification of model expressivity and theories of generalization. After reviewing this topic, I will discuss both recent applications with state-of-the-art results using novel tensor network architectures and recent progress in theories of machine learning based on unique capabilities of tensor networks. I will conclude by noting future areas of opportunities for research.

7.2	**Quantum exploration algorithms for multi-armed bandits** Daochen Wang, Xuchen You, Tongyang Li and Andrew Childs

Identifying the best arm of a multi-armed bandit is a central problem in bandit optimization. We study a quantum computational version of this problem with coherent oracle access to states encoding the reward probabilities of each arm as quantum amplitudes. Specifically, we show that we can find the best arm with fixed confidence using O(\sqrt{\sum_{i=2}^n \Delta_i^{-2}}) quantum queries (up to poly-logarithmic factors), where \Delta_i represents the difference between the mean reward of the best arm and the i-th best arm. This algorithm, based on variable-time amplitude amplification and estimation, gives a quadratic speedup compared to the best possible classical result. We also prove a matching quantum lower bound (up to poly-logarithmic factors).

7.3	**Learning Algorithms for Large-Scale Out-of-Equilibrium Quantum Systems** Eliska Greplova, Guliuxin Jin, Agnes Valenti, Julian Léonard and Sebastian D. Huber

Quantum simulators are at the forefront of large-scale quantum technologies [1, 2], where the state of the art system sizes can no longer be simulated classically. At the same time, precise verification and parameter estimation of these systems are necessary for further technological progress. Out-of-equilibrium dynamics present an avenue towards new unexplored territory [3] due to the challenges in both their theoretical description and experimental implementation. We introduce novel class of Hamiltonian learning [4] and parameter estimation [5] algorithms for out-of-equilibrium quantum simulations that can be scalably implemented on classical computers. We show how to use Bayesian learning and neural networks to reconstruct the physics of the large scale out-of-equilibrium quantum system from basic building blocs (as shown in Figure 1a). We illustrate the methods on simulated data and we argue that this class of methods can efficiently estimate all local parameters with previously inaccessible levels of precision (see Figure 1b). Additionally, we show that the amount of experimental data needed to achieve the desired precision is accessible using current experimental techniques. Our method constitutes a stepping stone toward scalable learning algorithms for large scale quantum experiments that can no longer be calibrated and verified using standard simulation methods.

7.4	**Learning quantum many-body states using generative models over the tensor network architecture** Khadijeh Najafi, Ahmadreza Azizi, Miles Stoudenmire and Masoud Mohseni

Generative models which learn the underlying probability distribution of the unlabeled data and generate new samples accordingly has become one of the cornerstones of probabilistic machine learning. Inspired by the probabilistic nature of quantum mechanics, we employ a generative model, known as the" Born machine" which uses quantum state representation and learns the joint probabilities over such quantum degrees of freedom. In this work, by leveraging on the power of the tensor networks for efficient expressivity, training, and sampling of a wide class of quantum many-body states, we investigate the trainability of two known classes of tensor network architectures known as matrix product state (MPS) and projected entangled pair state (PEPS) over quantum many-body states in one and two dimensions. In particular, we use the power of non-gradient descent-based learning algorithms that are adaptable with the tensor network architecture to learn various classes of many-body states.

8.1	**Promises and Challenges of Variational Quantum Algorithms** Patrick Coles

Beyond finding ground states, new applications have recently emerged for Variational Quantum Algorithms (VQAs). These include dynamical simulation, metrology, solving linear systems, and principal component analysis, among others. Moreover, new tools

are making VQAs more feasible, such as quantum-aware optimizers and error mitigation methods. VQAs also exhibit an intriguing resilience to certain types of noise. These signs point to VQAs soon providing quantum advantage for practically interesting problems. On the flip side, new analytical results have studied the gradient scaling for VQAs. These results paint a concerning picture, and show that barren plateaus in the training landscape will emerge if care is not taken. Barren plateaus are caused by lack of structure or simply by the presence of noise. Avoiding barren plateaus has now become a crucial area of research. This talk will review both the exciting and challenging nature of VQAs for near-term quantum computing.

8.2	**Quantum embeddings for machine learning** Aroosa Ijaz, Seth Lloyd, Maria Schuld, Josh Izaac and Nathan Killoran

Linear algebraic methods for machine learning, such as kernel methods, operate by embedding data points as vectors in a high dimensional vector space, and by applying methods of linear algebra to classify and to discriminate between embedded clusters of data. Quantum mechanics represents a natural setting for enacting such linear algebraic machine learning methods: data points can be embedded as quantum states (which are by definition vectors in high-dimensional vector spaces) and quantum computers can provide exponential speedups over classical computers for performing the linear algebraic techniques required to classify and to discriminate between those embedded states. This paper investigates the power of quantum embeddings for machine learning. We identify the optimal measurements to discriminate between clusters of data, and we show how to perform those measurements on existing quantum computers. We show that the performance of such quantum classifiers is completely determined by the quantum feature map that performs the embedding: we use adaptive learning techniques to find optimal quantum feature maps. Because of the phenomenon of quantum advantage, quantum embeddings are inaccessible to classical computers. Consequently, hybrid devices that combine classical preprocessing with quantum classification are strictly at least as powerful as classical kernel methods alone.

8.3	**Trainability of Quantum Neural Networks** Marco Cerezo, Akira Sone, Tyler Volkoff, Kunal Sharma, Lukasz Cincio and Patrick Coles

"Variational quantum-classical algorithms (VQCAs) optimize the parameters of a quantum neural network (QNN), V, to minimize a cost function, C. Many researchers believe that VQCAs will enable the first practical applications of noisy quantum computers. However, QNNs are heuristic methods with unproven scaling. In our works we rigorously study the trainability of two types of QNNs, the first is a parametrized quantum circuit commonly known as a layered hardware efficient ansatz, and the second is a dissipative perceptron-based QNN. For both of these QNNs our results state that choosing C to be a global function of V leads to an exponentially vanishing gradient (i.e., a barren plateau) even when V is shallow. This implies that many VQCAs proposed in the literature must revise their proposed cost functions. Second, we also show that under the same conditions, choosing C to be a

local function of V leads to a non-vanishing gradient, i.e., with the gradient vanishing no worse than polynomially. This suggests that QNNs have the potential to be trainable, if one chooses an appropriate cost function.

9.1	**Which functions can quantum machine learning models learn?** Maria Schuld

A lot of work in quantum machine learning focuses on how to practically train quantum models, or how to prove that they are classically intractable. However, an important question is which types of functions they can actually express. This talk presents insights from arXiv:2008.08605 that show how one can naturally write a quantum model as a partial Fourier series in the data. The accessible frequencies are determined by the nature of the data encoding gates in the circuit. This allows us to characterise the limits of what a quantum model can learn based on the tools of Fourier analysis - and it turns out that even classically intractable models can be trivial learners.

9.2	**Machine-learning tools for rapid control, calibration, and characterization of QPUs and other quantum devices** Nicolas Wittler, Federico Roy, Kevin Pack, Anurag Saha Roy, Max Werninghaus, Daniel Egger, Steffan Filipp, Frank Wilhelm and Shai Machnes

The current methodology of designing control pulses for superconducting circuits often results in an absurd situation: simplified analytic models which fail to predict gate fidelities to high accuracy, and calibrated pulse shapes which achieve good fidelities, but have no corresponding model. Moreover, for large number of qubits the QPU calibration is long and difficult, and determining a detailed error budget is nearly impossible. But without an error budget we have no insight into what is limiting QPU performance and how to improve it. To rectify the situation, we are implementing a novel procedure of Combined Calibration and Characterization (C3): An iterative combination of open-loop pulse optimization, model-free tune-up and iterative model fitting and refinement, utilizing a high-performance TensorFlow simulator. It will allow for a rapid, and largely automated bring-up process of QPUs. The result is a high-fidelity model, commensurate high-fidelity gates and a detailed error budget. The above components can then be utilized to implement machine-learning capabilities such as adversarial system characterization and automated experiment design, to further accelerate the process of gaining insight into the behavior of our systems. C3 software will be made available as an open-source project.

9.3	**Visualizing the loss landscape of quantum circuits** Kathleen Hamilton, Tyler Kharazi and Alex McCaskey

Variational training of quantum circuits is a hybrid quantum-classical workflow employed on near-term noisy intermediate scale quantum (NISQ) devices. In recent works [1,2] we have studied a specific example of variational training: data driven circuit learning (DDCL) using gradient-based optimization. With this workflow we have trained several circuits on superconducting qubits and explored how hardware noise affects training efficiency. In [1] we demonstrated that reducing two-qubit gates can

improve circuit trainability. In [2] we demonstrated that incorporating matrix-based error mitigation can also improve performance on noisy hardware. To further investigate the effect of noise in this hybrid workflow, the qualitative methods of loss landscape visualization introduced in [3,4] are applied to quantum circuits and the DDCL workflow. We are presenting results for two loss functions - the Jensen Shannon Divergence and the Maximum Mean Discrepancy. Multiple circuits on qubit registers up to 5 qubits and with up to 30 variational parameters are trained in noiseless simulation using stochastic gradient-based optimization (via Adam [5]). We use three methods to visualize the loss landscape. The first method sweeps out 1D traces to explore the structure of the minima around the final trained parameters. The second method uses a random sample of points in a 2D plane to visualize the topology of the landscape. The final method is the linear interpolation introduced in [3] applied to a random set of initial parameters and the final trained parameters. These methods are executed in noiseless simulation and also on superconducting qubit platforms. Understanding the loss landscape of quantum circuits and the effects of hardware noise will aid in developing efficient training methods. Barren plateaus in the loss landscape are a significant obstacle in gradient-based training of quantum circuits [6,7]. We observe relatively flat dimensions in the circuit gradient as well as instances of multiple local minima.

10.1	**Utilizing NISQ devices for evaluating quantum algorithms** Eleanor Rieffel

With the advent of quantum supremacy, we have an unprecedented opportunity to explore quantum algorithms in new ways. The emergence of general-purpose quantum processors opens up empirical exploration of quantum algorithms far beyond what has been possible to date. Challenging computational problems arising in the practical world are often tackled by heuristic algorithms. While heuristic algorithms work well in practice, by definition they have not been analytically proven to be the best approach or to outperform the best previous approach. Instead, heuristic algorithms are empirically tested on benchmark and real-world problems. With the empirical evaluation NISQ hardware enables, we expect a broadening of established applications of quantum computing. What to run and how best to utilize these still limited quantum devices to gain insight into quantum algorithms remain open research questions. We discuss opportunities and challenges for using NISQ devices to evaluate quantum algorithms, including in elucidating quantum mechanisms and their uses for quantum computational purposes, in the design of novel or refined quantum algorithms, in compilation, error-mitigation, and robust algorithms design, and in techniques for evaluating quantum algorithms empirically.

10.2	**Computing Edit Distances with D-Wave** Riccardo Mengoni, Francesco Trotti and Alessandra Di Pierro

Quantum Annealing (QA) can be used to approach computationally intensive aspects of Machine Learning (ML), provided that the algorithms are formulated as Quadratic Unconstrained Binary Optimisation (QUBO) problems. We consider the problem of calculating an edit distance between graphs and show a QUBO formulation of this

problem that makes it amenable for a QA treatment. Graph edit distances are measures of similarity between graphs that play important roles in several applications including kernel methods in ML. We show the results of testing our approach on the most advanced QA hardware available to the public, namely the D-Wave 2000Q.

10.3	**Quantum-tailored machine-learning characterization of quantum processors** Élie Genois, Agustin Di Paolo, Noah Stevenson, Gerwin Koolstra, Akel Hashim, Alexandre Blais and Jonathan A. Gross

Characterizing quantum devices by standard methods requires resources that grow exponentially with system size. This scaling arises from the need for acquiring complete information about a quantum state. Most of the time one already has some information about the device, however, and such rigorous characterization is unnecessary. Machine-learning (ML) provides an attractive framework for developing less demanding characterization and calibration protocols by processing limited amounts of quantum data heuristically. Here, we introduce a machine-learning architecture for inferring the dynamics of a quantum device from time-series measurement data. Our architecture is recurrent in nature and leverages quantum-mechanical structure in its design to interpret measurement data from quantum devices more efficiently. More precisely, we treat the case of a driven superconducting transmon qubit probed by a weak microwave tone implementing heterodyne measurement. We show how our custom, physics-inspired ML architecture can learn the dynamics of this system based on the measurement record and how this information can be used to infer the device parameters. We compare the performance of our custom ML model to off-the-shelf architectures based on recurrent neural-network units and to a physical model with parameters calibrated using standard experiments. We treat both numerically generated and experimental data to explore in detail the behavior of our model and demonstrate its utility in a real-world setting. This work shows how a physics-inspired ML model can improve the learning of device-specific quantum dynamics which can then be used to characterize and calibrate superconducting quantum processors.

11.1	**Simulating quantum systems with probabilistic models** Juan Carrasquilla

I will introduce a probabilistic representation of the quantum state in terms of the outcome distribution of a positive operator valued measure. In this language, time evolution translates into evolution of probability distributions subject to "somewhat" stochastic matrices, which are a generalization of stochastic matrices. I approximate the evolution of the quantum state using machine learning language models such as the Transformer and recurrent neural networks and provide a proof-of-principle demonstration of the approach on simple quantum circuits and on prototypical models in condensed matter and quantum information subject to dissipation.

11.2	**Error mitigation with Clifford quantum-circuit data** Piotr Czarnik, Andrew Arrasmith, Patrick Coles and Lukasz Cincio

Achieving near-term quantum advantage will require accurate estimation of quantum observables despite significant hardware noise. For this purpose, we propose a novel, scalable error-mitigation method that applies to gate-based quantum computers. The method generates training data $\{X_i^{\text{noisy}},X_i^{\text{exact}}\}$ via quantum circuits composed largely of Clifford gates, which can be efficiently simulated classically, where X_i^{noisy} and X_i^{exact} are noisy and noiseless observables respectively. Fitting a linear ansatz to this data then allows for the prediction of noise-free observables for arbitrary circuits. We analyze the performance of our method versus the number of qubits, circuit depth, and number of non-Clifford gates. We obtain an order-of-magnitude error reduction for a ground-state energy problem on 16 qubits in an IBMQ quantum computer and on a 64-qubit noisy simulator.

11.3	**Attention-based Quantum Tomography** Peter Cha, Paul Ginsparg, Felix Wu, Juan Carrasquilla, Peter McMahon and Eun-ah Kim

With rapid progress across platforms for quantum systems, the problem of many-body quantum state reconstruction for noisy quantum states becomes an important challenge. Recent works found promise in recasting the problem of quantum state reconstruction to learning the probability distribution of quantum state measurement vectors using generative neural network models. Here we propose the "Attention-based Quantum Tomography" (AQT), a quantum state reconstruction using an attention mechanism-based generative network that learns the mixed state density matrix of a noisy quantum state. The AQT is based on the model proposed in "Attention is all you need" by Vishwani et al (2017) that is designed to learn long-range correlations in natural language sentences and thereby outperform previous natural language processing models. We demonstrate not only that AQT outperforms earlier neural-network-based quantum state reconstruction on identical tasks but that AQT can accurately reconstruct the density matrix associated with a noisy quantum state experimentally realized in an IBMQ quantum computer. We speculate the success of the AQT stems from its ability to model quantum entanglement across the entire quantum system much as the attention model for natural language processing captures the correlations among words in a sentence.

12.1	**Entanglement Induced Barren Plateaus** Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe

In recent years the prospects of quantum machine learning and quantum deep neural network have gained notoriety in the scientific community. By combining ideas from quantum computing with machine learning methodology, quantum neural networks (QNNs) promise new ways to interpret classical and quantum data sets. In this work, we show that, with high probability, entanglement between the visible and hidden units can prevent learning. We connect quantum thermodynamics and quantum deep learning to demonstrate that the states on visible units are likely to be exponentially

close to a maximally mixed state. Moreover, the gradients of such models with respect to their parameters will vanish as the dimension of the system increases for reasonable random ensembles of quantum models. We call this effect entanglement induced barren plateaus and argue that their existence necessitates a re-evaluation of how we think about the role that entanglement and other quantum effects play within quantum machine learning algorithms.

12.2	**Entanglement and Optimization in Hamiltonian Variational Ansatz** Cunlu Zhou, Roeland Wiersema, Yvette De Sereville, Juan Carrasquilla, Yong-Baek Kim and Henry Yuen

Entanglement, the fundamental concept in quantum theory that describes non-local correlations in extended quantum systems, provides a unifying concept for the understanding of phenomena in a broad range of research areas related to quantum theory and quantum information science. Here, we use entanglement as a tool to study variational quantum circuits, their optimization dynamics, and their ability to approximate the ground state of prototypical many-body Hamiltonians in condensed matter physics. We focus on a special family of quantum circuits called the Hamiltonian Variational Ansatz (HVA), which takes inspiration from the quantum approximation optimization algorithm and adiabatic quantum computation. Through the study of its entanglement spectrum and energy gradient statistics, we find that HVA exhibits favourable structural properties such as mild or entirely absent barren plateaus and a restricted state space that ease their optimization in comparison to the commonly used ``hardware-efficient ansatz'' in the context of the variational quantum eigensolver (VQE) algorithm. We also numerically observe that the optimization landscape of HVA becomes almost trap free, i.e., no sub-optimal minima, when the ansatz is over-parameterized. We observe a size-dependent ``computational phase transition'' as the number of layers in the HVA circuit is increased where the optimization crosses over from a hard to an easy region in terms of the quality of the approximations and speed of convergence to a good solution. Unlike the analogous phase transitions observed in the learning of random unitaries where the transitions occurs at a number of layers that grows exponentially with the number of qubits, our numerical experiments suggest that the threshold in the number of layers in the circuit necessary to achieve the over-parameterization phenomenon scales linearly in the number of qubits. Lastly, as a demonstration of its entangling power and effectiveness, we show that HVA can find accurate approximations to the ground states of a modified Haldane-Shastry (MHS) Hamiltonian on a ring, which has long-range interactions and has a power-law entanglement scaling.

12.3	**Quantum enhancements for deep reinforcement learning in large spaces** Sofiene Jerbi, Lea Marion Trenkwalder, Hendrik Poulsen Nautrup, Hans Briegel and Vedran Dunjko

In the past decade, the field of quantum machine learning has witnessed an ever increasing attention due to its prospects of bringing genuine computational advantages to now widespread algorithmic methods. However, not all domains of

machine learning benefit equally from quantum enhancements. Notably, deep learning and reinforcement learning, despite their tremendous success both individually and combined, remain relatively unaddressed. One reason behind this is the appeal of the classical community for models and methods without prominent computational bottlenecks, leaving little room for quantum enhancements. In this work, we study the state-of-the-art neural-network approaches for reinforcement learning and demonstrate the substantial learning advantage that models with a sampling bottleneck can provide over conventional neural network architectures in complex learning environments. These so-called energy-based models, by requiring approximate sampling methods, trade off learning performance for efficiency of computation. Furthermore, we propose to leverage future and near term quantum algorithms to alleviate these computational costs. This is achieved using standard quantum computing machinery for the speed-up of classical methods and by employing generalized models to gain an additional quantum advantage.

13.1	**Operationally meaningful representations of physical systems in neural networks** Hendrik Poulsen-Nautrup

Recent years have seen a surge of interest in using machine learning techniques to make fundamental progress in science. Unlike humans however, current machine learning techniques lack the ability to describe physical systems in terms of representations endowed with a meaningful and abstract structure. Instead, most representation learning algorithms reflect statistical structure present in the training data and do not allow us to specify explicit and physically meaningful requirements on the representation. In my talk, I introduce a neural network architecture that learns representations of physical systems from experimental data such that different parameters in the representation are operationally meaningful within various experimental settings. I will provide an interpretation of this model in terms of communicating agents collectively learning a physical theory and present toy examples involving classical and quantum physics as well as the design of experimental settings.

13.2	**Improved direct state-overlap estimation for NISQ devices** Leonardo Guerini, Juan Carrasquilla and Leandro Aolita

We demonstrate a practical approach for simulating the SWAP test with no entangling gates. The method is based on a factorable quasi-probability representation of arbitrary states and fully-random single-qubit measurements, considerably outperforming previous proposals on several aspects.

13.3	**Fast adiabatic ground state preparation with few measurements** Benjamin F. Schiffer, Jordi Tura and J. Ignacio Cirac

Preparing the ground state of a Hamiltonian is a problem of great significance in physics with deep implications in combinatorial optimization. While the quantum adiabatic algorithm (QAA) [1] is known to return the ground state for sufficiently long preparation times, the decoherence times of noisy, intermediate-scale quantum

devices [2] render the QAA infeasible. In recent years, variational approaches such as the quantum approximate optimization algorithm (QAOA) [3] sparked intensive research interest, as they avoid long preparation times by relying on repeated measurements and a classical-quantum feedback loop [4]. However, optimization of a multiparameter cost function requires billions of measurements already for a moderate number of qubits, thus forsaking any practical quantum advantage. In our work, we aim at combining the strengths of the adiabatic and the variational approaches for fast and high-fidelity ground state preparation, benchmarking its results for large system sizes and quantum non-integrable local Hamiltonians.

14.1	**Small quantum computers and large classical data sets** Aram Harrow

Can a quantum computer help us analyze a large classical data set? Data stored classically cannot be queried in superposition, which rules out direct Grover searches, and it can often be classically accessed with some level of parallelism, which would negate the advantage of Grover even if it were possible. In this talk I will explore how to use quantum computers for data analysis tasks, such as maximum likelihood estimation, in the setting where the data set is too large to fit on the quantum computer, and at the same time, large classical computers are available.

14.2	**Learning quantum channels with tensor networks** Giacomo Torlai

We present a new approach to quantum process tomography, the reconstruction of an unknown quantum channel from measurement data. Specifically, we combine a tensor-network representation of the Choi matrix (a complete description of a quantum channel), with unsupervised machine learning of single-shot projective measurement data. We show numerical experiments for both unitary and noisy quantum circuits, for a number of qubits well beyond the reach of standard process tomography techniques.

15.1	**Quantum Sampling algorithm for Near-Term Devices** Dries Sels

Constructing and implementing useful quantum algorithms is one of the central challenges in quantum information science. Efficient sampling from a classical Gibbs distribution is an important computational problem with applications ranging from statistical physics over Monte Carlo and optimization algorithms to machine learning. Here, we introduce a family of quantum algorithms that provide unbiased samples by preparing a state that encodes the entire Gibbs distribution. We show that this approach leads to a speedup over a classical Markov chain for several examples including the Ising model and sampling from weighted, independent sets of two different graphs. We further propose a realistic implementation of sampling from independent sets based on Rydberg atom arrays. Our approach connects computational complexity with phase transitions, providing a physical interpretation

of quantum speedup, and opens the door to exploring potentially useful sampling algorithms using near-term quantum devices.

15.2	**Data re-uploading for a universal quantum classifier** Alba Cervera-Lierta

A single qubit provides sufficient computational capabilities to construct a universal quantum classifier when assisted with a classical subroutine. This fact may be surprising since a single qubit only offers a simple superposition of two states and single-qubit gates only make a rotation in the Bloch sphere. The key ingredient to circumvent these limitations is to allow for multiple data re-uploading. A quantum circuit can then be organized as a series of data re-uploading and single-qubit processing units. Furthermore, both data reuploading and measurements can accommodate multiple dimensions in the input and several categories in the output, to conform to a universal quantum classifier. The extension of this idea to several qubits enhances the efficiency of the strategy as entanglement expands the superpositions carried along with the classification. Extensive benchmarking on different examples of the single- and multi-qubit quantum classifier validates its ability to describe and classify complex data.

15.3	**Variational Quantum Linear Solver** Carlos Bravo-Prieto, Ryan LaRose, Marco Cerezo, Yigit Subasi, Lukasz Cincio and Patrick Coles

Previously proposed quantum algorithms for solving linear systems of equations cannot be implemented in the near term due to the required circuit depth. Here, we propose a hybrid quantum-classical algorithm, called Variational Quantum Linear Solver (VQLS), for solving linear systems on near-term quantum computers. VQLS seeks to variationally prepare $|x\rangle$ such that $A|x\rangle\propto|b\rangle$. We derive an operationally meaningful termination condition for VQLS that allows one to guarantee that a desired solution precision ϵ is achieved. Specifically, we prove that $C \geq \epsilon^2 / \kappa^2$, where C is the VQLS cost function and κ is the condition number of A. We present efficient quantum circuits to estimate C, while providing evidence for the classical hardness of its estimation. Using Rigetti's quantum computer, we successfully implement VQLS up to a problem size of 1024×1024. Finally, we numerically solve non-trivial problems of size up to $2^{50}\times2^{50}$. For the specific examples that we consider, we heuristically find that the time complexity of VQLS scales efficiently in ϵ, κ, and the system size N.