QTML 2018

Abstracts

The abstracts are sorted alphabetically with regards to the author's surname.
Send any queries or updates to schuld@ukzn.ac.za.

Invited Talks

A route towards quantum-enhanced artificial intelligence

Vedran Dunjko (LIACS, University of Leiden)

Artificial intelligence (AI) is a heavily overloaded term, which historically pertains to the development of human-level machine intelligence. What is meant by intelligence is always fuzzy, and leads to various notions of AI. These range from topics which involve very difficult concepts like ``consciousness'', to specialized endeavors like programming computers to play games, like chess, at or beyond human level. Another aspect of the latter type of "pragmatic AI" involves machine learning techniques, where machines learn to solve problems whose exact formal specification (unlike chess play) may be messy, but humans tend to do well. The majority of modern research in AI focuses on such more specific tasks, with emphases on machine learning aspects, and planning aspects (game play) — often relying on various search methodologies, and others. The latter often conceal NP-hard problems such as travelling salesman or boolean satisfiability. The field of quantum machine learning has provided a plethora of examples of quantum improvements in the machine learning aspects. However, quantum computing has also been proven useful for various (NP-hard) algorithmic problems, which are at the crux of the search aspects critical in AI. Does this mean that all the building blocks of "quantum AI" are already present? How much can quantum computing help?

In this talk we will reflect on aspects of quantum machine learning and of quantum-enhanced search algorithms — including some recent results showing even small quantum computers can help — specifically from the perspective of AI.

Gaussian boson sampling for molecular docking

Mark Fingerhuth (ProteinQure)

In this work, we show that Gaussian boson sampling (GBS) can be used for molecular docking in structure-based drug design. In one combinatorial formulation of the molecular docking problem, an isomorphous subgraph matching method was utilized to generate ligand orientations in the binding site. This method was demonstrated to be NP-complete, necessitating increasingly large computational resources as the ligand and binding site grows in complexity. For this reason, that we have developed an algorithm suitable for continuous variable quantum devices, which will enable molecular docking for large therapeutic macromolecules. We formulate the molecular docking problem as the problem of finding maximum vertex-weighted clique in the contact compatibility graph. Subsequently, we find that stochastic algorithms are enhanced through GBS, which selects dense subgraphs with high probability. By post-processing the resulting subgraphs we obtain solutions to the W-MAX-CLIQUE problem. These findings rely on a link between graph density and the number of perfect matchings – enumerated by the Hafnian – which is the relevant quantity determining sampling probabilities in GBS.

Machine Learning assisted Quantum Physics.

Patrick Huembeli (ICFO)

In the last few years Machine learning (ML) has been widely applied to quantum physics and has shown impressive and exciting results. In this talk I would like to take a step back and discuss about the usefulness of ML in physics. Especially in tasks where results are not certifiable and we are left to blindly trust the ML predictions. I will present two projects that show the potential of ML assisted research in physics where we could find an advantage over known methods. First, I will introduce the automatic discovery of characteristic features of phase transitions in many body localization (MBL), where machine learning helps to reduce the numerical effort of finding MBL phase transitions by orders of magnitude. Second I will present a quantum state tomography tool that uses a restricted Boltzmann machine as an ansatz which is scalable in system size and does not suffer from the exponential growth of the wavefunction. In both cases, the predictions of the learning algorithms are either certifiable or give probabilistic guarantees, and therefore they are useful in scientific inquiry. In the latter case, there is also a chance of using quantum sampling in the training phase of the algorithm, closing a loop between quantum-enhanced ML and ML used in quantum physics.

Quantum computation, security and machine learning: a look at adversarial quantum learning

Nana Liu (CQT University of Singapore)

The success of the modern internet relies in no small part on understanding the interplay between computation and security. More recently this area has also seen contributions from machine learning, including spam filters and malware detection. With the rising prevalence of machine learning algorithms, it is also important to address whether new security issues arise. Machine learning applications often require training or test data that originate from remote data centres or sensors. This decentralised set-up opens the door to adversaries that could exploit existing vulnerabilities in those algorithms. Evidence suggests that these vulnerabilities can grow with the dimensionality of the data. As quantum technologies become more accessible, these same concerns for quantum data is expected to become more important in a future quantum internet, especially as the most classically-intractable quantum systems of interest are usually high-dimensional. In this area which we can call adversarial quantum learning, we ask some basic questions: How do we verify high-dimensional quantum data with less resources? How do errors or adversarial changes in quantum data affect the probability of misclassification? Is there a dependence on dimensionality? How do we detect unusual quantum data? We take some first steps in addressing these questions and look towards the future of this intersection between computation and security on quantum data.

Machine Learning in Quantum Experiments

Hendrik Poulsen Nautrup (Innsbruck University)

Quantum experiments push the envelope of our understanding of fundamental concepts in quantum physics. However, further breakthroughs require more, and more complex problems to be solved and understood. Accordingly, designing more, and more complex experiments becomes difficult and often clashes with our "classical" human intuition. The question that I will address in my talk is whether machine learning can help where human intuition fails. I will introduce an autonomous learning model which learns to design complex photonic experiments, without relying on previous knowledge and often flawed intuition. Therefore, I will frame the design of quantum experiments in terms of a simple game which offers a pathway for machine learning techniques in experimental quantum physics. We will thoroughly explore the capacity of reinforcement learning in this game of quantum experiments: I will describe a system that can not only learn how to design novel experiments, but in the process also discovers nontrivial experimental techniques. The features of learning that we will come across during my talk support optimism with regard to a potentially more creative role of machine learning in research.

Applications of Quantum Autoencoders and Learning with Noise

Jonathan Olson (Zapata Computing)

A quantum autoencoder (QAE) is a NISQ quantum machine learning technique for learning to compress quantum data into a smaller subspace. In this talk, I will discuss some near-term applications of QAEs, including improving circuit ansatzes for state preparation algorithms like the variational quantum eigensolver. I will also share some observations on how QAEs can learn well in the presence of certain types of noise, and how other NISQ quantum machine learning algorithms might benefit.

A machine learning approach for benchmarking and training shallow quantum circuits

Alejandro Perdomo-Ortiz (Rigetti)

We present a data-driven quantum circuit learning (DDQCL) algorithm that can be used to assist the characterization of quantum devices and to train shallow circuits for generative models in machine learning. The procedure leverages quantum hardware capabilities to its fullest extent by using native gates and their qubit connectivity. We demonstrate that our hybrid approach can learn an optimal preparation of the maximally entangled Greenberger-Horne-Zeilinger states, and that it can efficiently prepare approximate representations of coherent thermal states; wave functions that encode Boltzmann probabilities in their amplitudes. Finally, complementing proposals to characterize the power or usefulness of near-term quantum devices, we provide a new hardware-independent metric called the qBAS score. It is based on the performance yield in a specific sampling task on one of the canonical machine learning data sets known as Bars and Stripes. We show how entanglement is a key ingredient in encoding the patterns of this data set into the quantum distribution resulting from the shallow quantum circuit; an ideal benchmark for testing hardware proposal starting at four qubits and up. We provide experimental of DDQCL on both, superconducting and ion-trap quantum computing platforms., highlighting the trade-off between several architectural circuit designs, including device connectivity and adjustable circuit depth.

Machine learning as a benchmark on current hardware

Raphael Poser (Oak Ridge National Laboratory)

The intersection of machine learning and quantum computing can yield interesting tests of quantum computer properties and important new metrics. One of the immediate applications on near term hardware include optimization of sampling from unknown probability distributions, which have applications in a broad array of quantum simulation, from chemistry to nuclear physics to field theory problems. Here we outline an extension to the quantum born machine sampling concept, which simultaneously implements training on quantum hardware while distilling a machine-learning-based benchmark. We will outline the relevant metrics and performance of superconducting hardware in the task of efficiently encoding data from a training set. In addition, we will discuss generative models for efficient ansatz state preparation circuits which maximize accuracy in simulation algorithms such as the variational quantum eigensolver.

Open quantum generalization of Hopfield neural networks

Pietro Rotondo (University of Nottingham)

I will propose a new framework to understand how quantum effects may impact on the dynamics of neural networks. I will show how to implement the dynamics of neural networks in terms of Markovian open quantum systems, which will allow to treat thermal and quantum coherent effects on the same footing. In particular, I will propose an open quantum generalization of the Hopfield neural network, the simplest toy model of associative memory. I will determine its phase diagram and show that quantum fluctuations give rise to a qualitatively new non-equilibrium phase, characterized by limit cycles (whose physical interpretation is unfortunately unclear).

In conclusion I will exhibit a quantum many-body system with disorder that displays the key features of an Hopfield-like memory and thus may guide upcoming experimental attempts to build quantum associative memories in lab.

Quantum Memristors and Neuromorphic Quantum Computing

Mikel Sanz (University of the Basque Country)

Technology based on memristors, resistors with memory whose resistance depends on the history charges crossing the device, has been a turning point to develop neuromorphic architectures for classical computation. However, in contrast to the known quantized models of passive circuit elements, such as inductors, capacitors or resistors, the design and realization of a quantum memristor was still missing. We introduce the concept of a quantum memristor as a quantum dissipative device whose decoherence mechanism is controlled by a continuous-measurement feedback scheme, which accounts for the memory. We provide numerical simulations showing that memory effects actually persist in the quantum regime. We show how to engineer the quasiparticle dynamics in Josephson junctions to construct quantum memristors in superconducting circuits and use our quantization method to design memristor-type constructions in other quantum platforms, particularly photonics. We use this building block to quantize the Hodgkin-Huxley model describing the propagation of the action potential in the axon of a neuron. Finally, we discuss the construction of quantum mem-perceptrons based on quantum memristors. Summarizing, quantum memristors might be considered, in the framework of quantum biomimetics, a building block for quantum neural networks, quantum machine learning, quantum simulations of non-Markovian systems, and in the long term, neuromorphic quantum computation.

Making quantum algorithms learn from data

Maria Schuld (University of KwaZulu-Natal, Xanadu)

An important question in the young discipline of quantum machine learning is what impact quantum computing technology will have on the field of machine learning. Can we accelerate known algorithms, or even contribute entirely new methods of data-driven decision making? How can in particular near-term devices - which run short and noisy quantum algorithms - be used to find patterns in data? One line of research, so called variational circuits, consider parameter-dependent quantum algorithms. These algorithms are interpreted as machine learning models that can be trained for a given task, for example to classify unseen data samples or to generate artificial data. In this talk I will give an overview of what we know and – more importantly – do not know about learning with variational circuits. I will cover ideas of how to train quantum algorithms, how to elegantly implement a neural network in a photonic quantum computer, and how we may potentially surpass purely classical models with this approach.

Machine learning for certification of photonic quantum information

Fabio Sciarrino (Sapienza Università di Roma)

Photonic technologies provide a promising platform to address at a fundamental level the connection between quantum information and machine learning. We will exploit machine learning as a tool to validate quantum devices such as Boson Samplers [1]. Indeed, the difficulty of validating large-scale quantum devices poses a major challenge for any research program that aims to show quantum advantages over classical hardware. To address this problem, we propose a novel data-driven approach wherein models are trained to identify common pathologies using supervised and unsupervised machine learning [2,3]. Our results provide evidence on the efficacy and feasibility of this approach, paving the way for its adoption in large-scale implementations.

[1] F. Flamini, N. Spagnolo, F. Sciarrino, “Photonic quantum information processing: a review”, [arXiv:1803.02790] *Reports on Progress in Physics*
[2] I. Agresti, N. Viggianiello, F. Flamini, N. Spagnolo, A. Crespi, R. Osellame06863], N. Wiebe, F. Sciarrino, “Pattern recognition techniques for Boson Sampling validation”, [arXiv:1712.06863]
[3] T. Giordani, F. Flamini, M. Pompili, N. Viggianiello, N. Spagnolo, A. Crespi, R. Osellame, N. Wiebe, M. Walschaers, A. Buchleitner, F. Sciarrino. “Experimental statistical signature of many-body quantum interference”. *Nature Photonics* **12**, 173 (2018).

Classical and Quantum Machine Learning with Tensor Networks

Miles Stoudenmire (Flatiron Institute)

In the last few decades, there has been great progress in algorithms for working with many-body quantum wavefunctions. Though most many-body wavefunctions are prohibitively large, interesting wavefunctions such as ground states of Hamiltonians have structures of entanglement which allow them to be compressed. The compression is carried out by representing the wavefunction as a tensor network, which are equivalent to certain interesting quantum circuits.

It turns out that tensor networks are actually a very general tool for compressing large tensors. I will discuss an interesting framework for machine learning where large tensors naturally arise and can be successfully represented by tensor networks of the same type used in quantum physics. The approach has many interesting payoffs for machine learning, such as significantly better scaling for kernel-learning models and adaptive training algorithms.

Perhaps the most interesting direction for using tensor networks for machine learning is that the same models one uses classically can be straightforwardly evaluated on quantum hardware. Models can be pre-trained classically then extended to use quantum resources. I will discuss proof-of-principle numerical experiments, including robustness to qubit noise, and conclude by discussing the feasibility of evaluating models for realistic data sets on small, near-term quantum computers.

Contributed Talks

Exact compression of quantum states and effective ansatz for quantum autoencoder

Yudong Cao (Zapata Computing)

A major advantage of the variational quantum algorithm paradigm is that it makes substantial use of both quantum and classical processors. This hybrid approach is particularly appealing in the current era of noisy intermediate scale quantum (NISQ) devices where qubits are expensive and sensitive to error and noise. However, one potential challenge in implementing variational quantum algorithms effectively concerns optimizing the quantum circuit. As the problem size grows, the number of parameters that need to be optimized also grows, giving rise to high-dimensional black-box optimization problems that are difficult to tackle for a classial computer. Such difficulty is rigorously manifested in a recent study showing that the optimization landscape has vanishingly small gradient in most regions. To surmount the problem, one potential solution is to start from an educated initial guess of a good quantum circuit for the problem, a circuit ansatz so to speak, instead of a quantum circuit chosen in a random or ad hoc manner. As an example, for variational quantum eigensolvers (VQE) there are excellent quantum circuit ansatz constructions available for approximating the ground state of a physical system. However, for quantum autoencoders it remains unclear which circuit ansatzes may be useful.

The technique we develop here addresses this issue by approximating the input quantum states with matrix product states. For this approximate representation of the data, then, we can exactly compute a compression scheme. Although in practice the quantum states that need to be compressed may be more complicated than their MPS approximations, it is nonetheless useful to first compress instead approximate representation of the quantum states and use the compression circuit as an educated initial guess, or an ansatz, for further training a quantum circuit to compress the original set of quantum states. Specifically, one may introduce additional tuning parameters to the circuit ansatz and fine tune the parameters to account for the error incurred in the state approximations used for constructing the ansatz.

Optimal universal learning machines for quantum state discrimination

Marco Fanizza (Scuola Normale Superiore di Pisa)

A basic machine learning setting is supervised learning, which deals with the task of inferring a labeling rule on a data set, given a certain number of labeled points.
As a quantum generalization of supervised learning, we consider the problem of correctly classifying a given quantum two-level system (qubit) which is known to be in one of two equally probable quantum states. We assume that this task should be performed by a quantum machine which does not have at its disposal a complete classical description of the two template states, but can only have partial prior information about their level of purity and mutual distance. Moreover, similarly to the classical supervised learning paradigm, we assume that the machine can be trained by n qubits prepared in the first template state and by n more qubits prepared in the second template state. In this situation we are interested in the optimal process which correctly classifies the input qubit with the largest probability allowed by quantum mechanics. The problem is studied in full generality for a number of different prior information scenarios and for an arbitrary size n of the training data. Finite size corrections around the asymptotic limit $n\rightarrow \infty$ are derived.

Witnessing non-Markovianity with quantum memory

Christina Giarmatzi (University of Queensland)

A challenge present in every experimental setup is to monitor the interaction between some desired system and their environment. A typical setup has various elements that perform the channels between measurements. Ideally, the channels are the same in each run of the experiment. If they are not the same, it could be due to external factors, like the temperature of the lab, but it could also be due to different states of the initial system. We are interested in the latter case and the type of correlations that can exist between the initial state and the channels. The process is called Markovian when in every run of the experiment the channels are independent of the initial state of the system. When the process is non-Markovian, it can be simulated with an extra memory to carry the extra correlations which can be quantum or classical. We define “non-Markovianity with classical” or “with quantum memory” depending on the type of memory required to simulate the process. In non-Markovian processes with classical memory, the channels are essentially controlled classically by the input state. For non-Markovian processes with quantum memory there is another quantum system that gets entangled with our system and carries the extra correlations. We present a way to distinguish between the two types of non-Markovianity by mapping our problem to the separability problem. Hence we can use all the known techniques to detect entanglement which maps to detecting non-Markovianity with quantum memory. Finally, to illustrate our method we use two different families of separability criteria to obtain witnesses, one of them required convex optimisation methods (SemiDefinite Programming). As an example of a non-Markovian process with quantum memory we used a setup where the environment and the system interact with the Hamiltonian according to the Heisenberg model. This distinction of the two types of non-Markovianity is important in every experiment setup when looking for the nature and source of noise. Our development of witnesses enables us to investigate the nature of noise without the need of full process tomography and paves the way for compensation techniques.

Neural-network and tensor-network duality: applications in quantum physics and machine learning

Ivan Glasser (Max Planck Institute of Quantum Optics)

Learning a probability distribution from data and approximating a quantum state are two fundamentally related problems. Techniques designed to represent probability distributions, such as Boltzmann machines, have been applied as ansatz for a many-body wave function, while tensor networks designed to represent quantum states have been used in machine learning. We discuss how these distinct approaches are related, with a focus on the relationship between probabilistic graphical models and tensor networks. We show that the two frameworks can be connected through the definition of generalized tensor networks

which allow tensors to be copied. This correspondence has applications for particular models in both quantum physics and machine learning.

On the one hand it implies that restricted Boltzmann machines are related to previously introduced generalized tensor networks such as String-Bond States. These provide a natural generalization of restricted Boltzmann machines which is also applicable to quantum states with larger local Hilbert space. This connection sheds light on the underlying architecture of restricted Boltzmann machines and their efficiency at representing many-body quantum states. We provide ways of combining these classes of states together and benchmark these techniques on the challenging problem of representing chiral topological states, giving analytical as well as numerical evidence that these networks are able to approximate a chiral spin liquid with high accuracy.

On the other hand these generalized tensor networks can also be used in machine learning. They overcome some of the limitations of regular tensor networks in higher dimensions, while keeping the computation efficient. We provide an algorithm to train these networks in a supervised learning context as well as a method to combine neural networks and tensor networks together. We benchmark our algorithm for several generalized tensor-network architectures on the task of classifying images and sounds, and show that accurate classification can be obtained while keeping the bond dimension small. Our models correspond to functions that can be realized by quantum circuits and may guide the development of near-term quantum machine learning architectures. In particular our results show that quantum circuits might be able to perform accurate classification, but that good regularization techniques also need to be developed.

- Ivan Glasser, Nicola Pancotti, Moritz August, Ivan D. Rodriguez and J. Ignacio Cirac, 'Neural-Network Quantum States, String-Bond States, and Chiral Topological States', Physical Review X 8, 011006 (2018)
- Ivan Glasser, Nicola Pancotti and J. Ignacio Cirac, 'Supervised learning with generalized tensor networks', arXiv:1806.05964

Problem Solving Optimization by Machine Learning for Gate-Model Quantum Computers

Laszlo Gyongyosi (University of Southampton)

Quantum computers exploit the fundamentals of quantum mechanics to solve computational problems more efficiently than traditional computers [1-2]. Gate-model quantum computers provide a flexible framework to realize quantum computers in experiments [3-6]. The maximization of the objective function of computational problems is a remarkable application scenario of experimental gate-model quantum computers. The objective function estimation of the quantum computer is a high-cost procedure that requires several rounds of quantum state preparations, quantum computational steps, and quantum state measurements. Here, we define a machine-learning-based framework for objective function estimation and maximization in gate-model quantum computers. The method significantly reduces the costs of the objective function estimation and provides an estimate of the new state of the quantum computer. The framework integrates an objective function extension procedure, a segmentation algorithm that utilizes the gate parameters of the quantum computer, and a machine-learning unit for the quantum state prediction. The results are particularly convenient for the performance optimization of experimental gate-model quantum computations.

[1] Biamonte, J. et al. Quantum Machine Learning, Nature, 549, 195-202 (2017).
[2] Farhi, E., Goldstone, J., Gutmann, S. and Neven, H. Quantum Algorithms for Fixed Qubit Architectures, arXiv:1703.06199v1, (2017).
[3] Lloyd, S., Mohseni, M. and Rebentrost, P. Quantum algorithms for supervised and unsupervised machine learning, arXiv:1307.0411v2, (2013).
[4] Lloyd, S., Mohseni, M. and Rebentrost, P. Quantum principal component analysis, Nature Physics, 10, 631 (2014).
[5] L. Gyongyosi, S. Imre and H. V. Nguyen, A Survey on Quantum Channel Capacities, IEEE
Communications Surveys and Tutorials 99, 1, doi: 10.1109/COMST.2017.2786748 (2018).
[6] S. Imre and L. Gyongyosi. Advanced Quantum Communications - An Engineering Approach. Wiley-IEEE Press (New Jersey, USA), (2013).

Discovering physical concepts with neural networks

Raban Iten (ETH Zurich)

The formalism of quantum physics is built upon that of classical mechanics. In principle, considering only experimental data without prior knowledge could lead to an alternative quantum formalism without conceptual issues like the measurement problem. As a first step towards finding such an alternative, we introduce a neural network architecture that models the physical reasoning process and can be used to extract simple physical concepts from experimental data in an unbiased way. We apply the neural network to a variety of simple physical systems in classical and quantum mechanics, like damped pendulums, two-particle collisions, and qubits. The network finds the physically relevant parameters, exploits conservation laws to make predictions, and can be used to gain conceptual insights. For example, given a time series of the positions of the Sun and Mars as observed from Earth, the network discovers the heliocentric model of the solar system - that is, it encodes the data into the angles of the two planets as seen from the Sun.

Quantum optimization using the Gradient method

Thomas Konrad (University of KwaZulu-Natal)

We discuss a quantum algorithm to find a minimum of a function using a gradient descent combined with a Grover search. We show that this algorithm is faster than a simple classical steepest descent and faster than Grover's search algorithm for convex functions of more than two real arguments.

High-fidelity conditional two-qubit swapping gate

Niels Jakob Søe Loft (Aarhus University)

The promising nature of scalable quantum computing relies crucially on high-fidelity entangling operations. I will demonstrate that four coupled qubits can operate as a high-fidelity two-qubit entangling gate that we call the ZSWAP. The gate operation is tuned on and off controlled by the state of two ancilla (control) qubits. By using realistic device and noise parameters from state-of-the-art superconducting qubits, I will show that the conditional ZSWAP operation can be implemented with a fidelity above 0.99 in a time around 60 ns.

Homological analysis of multi-qubit entanglement

Riccardo Mengoni (University of Verona)

We propose the usage of a Topological Data Analysis tool namely Persistent Homology to characterize multipartite entanglement. On a multi-qubit data set we introduce metric-like measures defined in terms of bipartite entanglement and then we derive barcodes. We show that, depending on the distance, they are able to produce different classifications.

In one case, it is possible to obtain the standard separability classes. In the other case, a new classification of entangled states of three and four qubits is provided.

Duality Quantum Neural Networks

Wilson de Oliveira (Universidade Federal Rural de Pernambuco)

We present and investigate a novel generalised model of a quantum neural networks. This model is based on a new kind of quantum computer named duality quantum computer (DCQ) \cite{long2011duality}. We show that the this neural model subsumes both classical weighted and weightless neural networks as well their quantum counterparts. Can be seen as the quantum version of the vector space neural networks \cite{de2014vector}. We indicate how training can be effected and give directions for further directions particularly for solving NP-complete problems in polynomial time.

Supported by the Brazilian research council (grant CNPq-Universal 421849/2016-9) and by the Serrapilheira Institute (grant number Serra-1709-22626)

Quantum-Classical Reinforcement Parity Learning from Noisy Classical Data

June-Koo Kevin Rhee (Korea Advanced Institute of Science and Technology)

Recent discoveries of quantum algorithms for data analysis and machine learning have garnered much attention, and stimulated further exploration of quantum technologies that can provide considerable advantages to practical applications in the big data era. On the other hand, any physical implementations are subject to errors, and the quantum advantage must be retained in the realistically noisy settings. One interesting example in which the quantum algorithm outperforms the classical counterpart in the presence of noise is the problem of learning a hidden parity with noise (LPN).
In the classical parity learning problem, an oracle generates a uniformly random input $x\in\lbrace 0,1\rbrace^n$ and computes a Boolean function f_s defined by a hidden bit string $s\in\lbrace 0,1\rbrace^n$, $f_s(x)=\sum_{j=1}^n s_jx_j\;\text{mod }2$, where $s_j\;(x_j)$ is the jth bit of $s\;(x)$. A query to the oracle returns $\left(x,\;f_s(x)\right)$, and a learner tries to reconstruct s by repeating the query. An example of such application is low-density parity-check (LDPC) coding used in a noisy channel. In the presence of noise, the learner obtains $f'_s(x)$, where the prime $'$ indicates a bit flip error occurs with a certain probability.
In the quantum setting, the learner has access to a quantum oracle which implements a unitary transformation on the computational basis states and returns the equal superposition of $|x\rangle|f_s(x)\rangle$ for all possible inputs of x. Then the quantum learner with the ability to apply the Hadamard gate to all qubits before the readout can solve the problem efficiently. Learning from noiseless oracle is easy to both classical and quantum learners. However, in the presence of noise, only the quantum approach remains to be efficient. Therefore, the LPN problem serves as an exciting example for which the quantum technology can provide substantial enhancement even under inevitable imperfections.
However, in practical circumstances, a learner is most likely to receive noisy classical data, rather than noisy quantum data as considered in the quantum LPN algorithm. Moreover, the input x also experiences noise. Then, whether the quantum technology can still be beneficial is an interesting open problem.
In this work, we develop a quantum-classical hybrid algorithm for reinforcement learning of the hidden parity from the noisy classical data. A reinforcement learning approach to the problem can be done in the following steps. First, collect M pairs of the noisy classical data, $\left(x',\;f''_s(x)\right)$ for training, where prime symbols $'$ and $''$ represent faulty outcomes due to independent bit flip errors, and use an efficient means, such as quantum random access memory (QRAM), to convert the classical data to a quantum superposition

state $\sum_{j}^{M} |x_j' \rangle|f''_s(x_j)\rangle$. Then the quantum learning for M training data produces a measure of reward by applying the Hadamard gate, similar to the original quantum LPN algorithm. Here, the reward is given by the probability p_k assigned to candidate strings $\tilde{s}_k$ as the outcomes of the LPN algorithm. Finally for each training round, one can choose an action, $s^* = \tilde{s}_k$ for the maximum p_k, to prepare the LPN input quantum state for the next round. The iteration continues with acquiring the next training data set and combining it with the existing quantum state using another QRAM process.

We numerically test various strategies to choose an action s^* for updating the input state to LPN, and demonstrate significant improvement in the query complexity with the quantum-classical hybrid reinforcement learning compared to the direct application of the original quantum parity learning algorithm to classical data.

Learning Quantum Parity Oracle With Maximally Mixed Input Qubits

Daniel Kyungdeock Park (Korea Advanced Institute of Science and Technology)

Experimental realizations of quantum information processing have made remarkable progresses in the past decade. However, universal fault-tolerant quantum computation is still far from within reach. While efforts towards developing full-fledged quantum computers continue, identifying well-defined computational tasks for which weaker but more realistic devices can exceed classical counterparts is imperative.

Machine learning is an interesting family of problems for which near-term quantum devices can provide considerable advantages. In particular, the exponential quantum speedup is possible in the problem of learning a hidden parity function defined by an unknown binary string in the presence of noise (LPN). In the quantum LPN algorithm, all possible input binary strings are encoded in the data qubits as the equal superposition state, and a query to the quantum oracle stores the parity outcomes in the result qubit. Then the quantum learner with the ability to coherently rotate all qubits before the readout can solve the problem efficiently. However, the number of queries increases as the depolarizing rate increases, and the parity function can only be guessed with a success probability decreasing exponentially with the problem size if the data qubits are fully depolarized. Classically, the corresponding task can only be solved via brute-force enumeration in an exponentially large search space, provided that an efficient means to verify the answer exists.

Recently, we showed that the LPN problem with all input qubits maximally mixed except for the result qubit fits naturally in deterministic quantum computation with one qubit (DQC1). In principle, the DQC1-based LPN algorithm (DQC1-LPN) can reveal the hidden parity function efficiently via calculating the trace of the unitary operator that implements the quantum oracle. However, in practice, the performance of DQC1-LPN can degrade due to the noisy expectation value measurement for the trace estimation, and the number of measurements can increase rapidly as the length of the hidden bit string increases for a fixed accuracy.

The DQC1 model can be viewed as a special instance of the famous quantum phase estimation (QPE) algorithm. Here, we show that the QPE can be used to count the number of degenerate eigenvalues of a unitary operator acting on the maximally mixed state with the number of queries independent of the size of the unitary matrix. By exploiting this result, we show that the quantum LPN problem with fully depolarized data qubits can be solved efficiently using the QPE algorithm if another qubit is added to the DQC1 model. Unlike the DQC1-LPN algorithm, the number of measurements in the QPE-LPN algorithm for a desired accuracy is independent of the problem size. Hence, surprisingly, the addition of one qubit significantly enhances the learning speed.

Our work serves as an intriguing example illustrating that a classically intractable learning task can be formulated as the trace estimation or counting the degenerate eigenvalues that can be solved efficiently using only one or two qubits, and that the power of two qubits is significantly greater than that of one.

Hybrid quantum-classical schemes for generative adversarial learning: HQGANs

Jonathan Romero (Harvard University)

Quantum computing and machine learning are two fast-growing research areas that are improving our ability to process information and to understand and manipulate data, respectively. Recently, these two areas have been merged into the field of quantum machine learning (QML), seeking to find ways in which quantum computers can offer advantages at solving machine learning problems over classical computers. Here, we propose to use quantum computers to learn models that mimic observed data distributions, a type of task known as generative learning, by substituting neural networks with variational quantum circuits in the generative adversarial networks (GANs) framework. GANs are statistical models that learn to generate samples from an observed data distribution by looking at individual samples. They consist of two neural networks, known as the discriminator and the generator, competing against each other in a minimax game. This approach has proven very successful at modeling unknown data distributions for different applications, including image synthesis, semantic image editing, and more recently on molecular and materials discovery.

While a version of quantum GANs that encodes information as quantum states have been recently proposed, a quantum GAN than can be implemented on near-term quantum devices and that can be applied to classical data has not been realized. We propose a hybrid-quantum classical scheme for GANs, where variational quantum circuits are employed as generator and discriminator models. In this proposal, we outline the general theoretical framework of the HQGANs model and its implementation. Specifically, we address the following challenges: (1) implementing methods for encoding classical data on quantum registers, (2) developing quantum circuit architectures for discriminative and generative models, (3) designing suitable cost functions and optimization approaches and (4) devising strategies for implementing the model on NISQ devices. The proposed scheme might benefit machine learning by improving the ability to model more complex data distributions and could offer a new niche of applications for near-term quantum computers.

Searching for Majorana Zero Modes Using Model-free Reinforcement Learning

Makhamisa Senekane (National University of Lesotho)

Majorana fermions are particles which are their own antiparticles; hence they have zero charge. They are governed by non-Abelian statistics. For a Majorana fermionic operator γ, and the Hamiltonian of a system H, Majorana fermions satisfy fermionic anti-commutation relation and a Majorana fermion squares to 1. If, in addition to this, the fermionic operator commutes with the Hamiltonian of the system, then such an operator is a Majorana zero mode (MZM). Majorana zero modes are Majorana fermions bound to zero energy. MZMs have applications in both topological quantum computation and spintronics. In this work, we report the algorithm that searches for MZMs using reinforcement learning. Reinforcement learning is a machine learning paradigm where the learner is a decision-maker (agent) that takes action in an environment and receives rewards or penalties for the actions taken. Results obtained from this work demonstrate the significance of using reinforcement learning in the quest for Majorana zero modes.

Unsupervised classification of quantum data

Gael Sentís (University of Siegen)

Unsupervised learning algorithms are set to identify structure in unlabeled data, and they represent an extremely important class of algorithms in classical machine learning, due to the abundance of unlabeled over labeled data. In the context of a classification problem, a clustering algorithm assigns different labels to unlabeled data points that presumably come from sampling different probability distributions. In a scenario where data points are not vectors of classical variables but unknown states of quantum particles, the principles behind classical unsupervised learning algorithms, based on performing repeated operations over each data point, immediately fail to apply due to the impossibility of cloning quantum information. At the same time, new genuinely quantum strategies are possible, like reading out a collective property of sets of quantum particles by measuring them coherently. Thus, single-shot unsupervised quantum learning algorithms that deal directly with unlabeled information in a quantum form will necessarily have radically different capabilities and structure.

We investigate these questions by approaching the simplest case of unsupervised binary classification of quantum states: we are given a string of N quantum particles such that each of them can be in one of two unknown possible states, in any possible order; then, we are set to design the quantum algorithm that is able to best guess simultaneously how many particles are of each type and in which order. Since the two classes of states are unknown, the only accessible information to the observer resides in the correlations between the particles. We address this problem as a state discrimination task between 2^N hypotheses. We derive the optimal quantum algorithm that solves this task, based on a global measurement of the N particles, and obtain analytical formulae for its performance according to different figures of merit. These figures set ultimate performance bounds on any conceivable algorithm for unsupervised classification of quantum data. We then use these results to examine the existence of a performance gap with respect to online quantum algorithms, that operate individually on each particle, and with the classical analogue of this problem where quantum states are replaced by classical biased coins.

Quantum enhanced cross-validation with a parametric probabilistic quantum memory

Adenilton da Silva (Universidade Federal Rural de Pernambuco)

We show that probabilistic quantum memories can be used to speed up cross-validation. To achieve this task, we show that PQM can run in actual quantum computers, define a parametric quantum probabilistic memory and perform a classical experiment with neural networks to verify that the quantum enhanced cross-validation can be used to perform model selection.

This work was supported by the Serrapilheira Institute (grant number Serra- 1709-22626)

Reinforcement Learning Methods for Quantum Error Correction

Ryan Sweke (Freie Universitaet Berlin)

To date, the primary obstacle preventing the implementation of a scalable universal quantum computer remains the errors which are necessarily introduced when realizing physical qubits and implementing gates. The most promising and experimentally pursued strategy for overcoming this obstacle involves redundant encoding of quantum information in topological error correction codes. This strategy requires active error correction, via classical algorithms which are capable of diagnosing the underlying errors from syndrome measurements, and then suggesting appropriate corrections - a process known as decoding. While various decoding algorithms already exist - most prominently minimum weight perfect matching - in practice there are various challenges still remaining. In particular, the run time of the algorithms and the implicit assumptions made on the underlying noise model restrict their usefulness within currently available devices. We present decoding algorithms, obtained via various reinforcement learning techniques, for performing decoding within a variety of physically motivated contexts, and explore the extent to which these algorithms and techniques mitigate the disadvantages inherent in alternative decoding algorithms. In particular, we examine decoding algorithms obtained via deep-Q learning, as well as via Monte-Carlo tree search based approaches, as used in AlphaZero. Our work distinguishes itself from previous machine learning based decoders in two ways: Firstly, our approach allows us to obtain decoding algorithms capable of performing decoding at intermediate stages of quantum computation in the presence of faulty syndrome measurements, thus providing an experimentally relevant tool for fault tolerant quantum computation. Secondly, we provide the first decoding algorithms obtained via reinforcement learning, and illustrate the flexibility and suitably of this approach for this problem.

Deep reinforcement learning for quantum memory

Petru Tighineanu (Max Planck Institute for the Science of Light)

We employ reinforcement learning to train an artificial neural network that discovers, tabula rasa (i.e., with no human knowledge or guidance), complete quantum-error-correction strategies in a collection of quantum bits subject to decoherence [1]. The network discovers optimal encoding and decoding protocols of the logical qubit as well as intricate correction strategies that are adapted to the measurement outcomes. A key novelty of our work is the development of an immediate reward scheme based on a physically meaningful quantity describing the capacity to protect the quantum information stored in the quantum memory. Our work opens the prospect for developing fully automated quantum-error-correction strategies in complex quantum systems.

In the first part of the talk I will explain how the neural network develops on its own the building blocks of a quantum-error-correction strategy: the encoding of the logical qubit, the error detection and correction, and the decoding of the logical qubit at the end of the simulation time. In the second part I will explain how our approach can be used to develop strategies that are fully adapted to the subtleties of the quantum device such as the hardware specifications and the noise. I will also present a detailed analysis of the internal workings of the neural network.

[1] T. Fosel, P. Tighineanu, T. Weiss, and F. Marquardt, https://arxiv.org/abs/1802.05267, accepted in PRX https://journals.aps.org/prx/accepted/84071KecE0a11c0af34d321704c757b6149499a79

Experimentally Detecting a Quantum Change Point via Bayesian Inference

Shang Yu (University of Science and Technology China)

Detecting a change point is a crucial task in statistics that has been recently extended to the quantum realm. A source state generator that emits a series of single photons in a default state suffers an alteration at some point and starts to emit photons in a mutated state. The problem consists in identifying the point where the change took place. In this work, we build a pseudo-on-demand single photon source to prepare the photon sequences, and consider a learning agent that applies Bayesian inference on experimental data to solve this problem. This learning machine adjusts the measurement over each photon according to the past experimental results finds the change position in an online fashion. Our results show that the local-detection success probability can be largely improved by using such a machine learning technique. This protocol provides a tool for improvement in many applications where a sequence of identical quantum states is required.

Machine Learning in Quantum Experiments

Hendrik Poulsen-Nautrup

Quantum experiments push the envelope of our understanding of fundamental concepts in quantum physics. However, further breakthroughs require more, and more complex problems to be solved and understood. Accordingly, designing more, and more complex experiments becomes difficult and often clashes with our "classical" human intuition. The question that I will address in my talk is whether machine learning can help where human intuition fails. I will introduce an autonomous learning model which learns to design complex photonic experiments, without relying on previous knowledge and often flawed intuition. Therefore, I will frame the design of quantum experiments in terms of a simple game which offers a pathway for machine learning techniques in experimental quantum physics. We will thoroughly explore the capacity of reinforcement learning in this game of quantum experiments: I will describe a system that can not only learn how to design novel experiments, but in the process also discovers nontrivial experimental techniques. The features of learning that we will come across during my talk support optimism with regard to a potentially more creative role of machine learning in research.

Tutorials

Open-Source Software for Quantum Computing

Mark Fingerhuth (ProteinQure)

Quantum Machine Learning with the IBM Q

Part A - Introduction Waheeda Saib, Ismail Akhalwaya (IBM Africa Lab)

Part B - Carsten Blank (Data Cybernetics)

Quantum Machine Learning with Strawberry Fields

Maria Schuld (University of KwaZulu-Natal, Xanadu)

Posters

The posters will be displayed in the conference room throughout the week. On the first Tuesday coffee break we will have a 'poster presentation' where the authors are available at their posters.

Application of the machine learning techniques to the quantum optical problems.

Anesan Reddy, Ilya Sinayskiy, Francesco Petruccione

Due to the tremendous progress in the development of Graphics Processing Unit (GPU) technology we see the real-world application of machine learning techniques for various tasks (face recognition, natural language processing, etc). Recently, researchers have begun to apply machine learning Techniques to hard problems in physics. For example, Convolutional Neural Networks (CNN) are used to detect the phase transitions in Ising and Bose-Hubbard models. In this poster, I will report on the application machine learning techniques to the description of the dynamic and steady-state properties of the fundamental model describing light-matter interaction (dissipative Jaynes–Cummings model).

Learning nonlinear input-output maps with dissipative quantum systems

Jiayin Chen, Hendra Nurdin

Currently, there is a substantial interest in implementing hybrid quantum-classical algorithms on near-term noisy intermediate quantum technology (NISQ) for temporal learning problems. We develop a general theory for learning arbitrary nonlinear input-output maps with fading memory using dissipative quantum systems, as a theoretical foundation for temporal learning with hybrid processors. The theory identifies conditions required for a class of dissipative quantum systems to be an universal approximator for input-output maps with fading memory. We also introduce a concrete class of dissipative quantum systems and prove its universality. Numerical experiments demonstrate that for certain benchmark learning tasks, this universal class of quantum dissipative systems with a small number of qubits compares favorably to classical learning schemes such as the echo state networks and the Volterra series with a large number of programmable parameters.