



**QUANTUM MACHINE LEARNING MEETINGS
KWAZULU-NATAL, SOUTH AFRICA**

**Quantum Machine Learning Workshop 18-22 July 2016
Palm Beach Lodge, North Coast**

Abstracts

Mikhail Altaisky
Space Research Institute RAS

On the perspective of implementation of quantum neural networks using quantum dots

We propose an implementation of quantum neural networks using an array of quantum dots with dipole-dipole interactions. We demonstrate that this implementation is both feasible and versatile by studying it within the framework of GaAs based quantum dot qubits coupled to a reservoir of acoustic phonons. Using numerically exact Feynman integral calculations, we have found that the quantum coherence in our neural networks survive for over a hundred ps even at liquid nitrogen temperatures (77 K), which is three orders of magnitude higher than current implementations which are based on SQUID-based systems operating at temperatures in the mK range. The entanglement properties of such systems are discussed.

Ref: Applied Physics Letters 108 (2016) 103108.

**Elizabeth Behrman
Wichita State University**

Programming Quantum Hardware with Machine Learning: Quantum Annealing

Our goal is to devise a systematic method, using machine learning, to “program” a large-scale quantum computer to anneal to a desired state. Quantum computation has enormous potential, for which entanglement is key. A great deal of effort in recent years has produced methods to create GHZ states fairly reliably over a few qubits; our methods show promise to making it possible to create GHZ states over hundreds of qubits but also to tailor the particular entanglement desired for a particular computation, for example, a W state and others in addition to the GHZ state.

In our prior work we have used machine learning, via quantum learning algorithms, to “program” a simulated computer built with quantum bits (qubits) to do a variety of proof of concept tasks: Classical logic gates, quantum logic gates, associative memory tasks, computing an entanglement witness, and phase correction of entanglement estimates. For these quantum computers, the inputs are the initial state of the system and the outputs are a measure of the final state of the system. These methods use machine learning to determine, via supervised learning, the physical parameters of the qubits, and their pairwise coupling strengths, that are the “program” that causes the system to calculate the correct paired outputs for a given set of input data.

Here we apply machine learning to “program” a quantum annealing computer, originally designed to solve binary optimization problems. In simulation, it has been programmed to anneal to the GHZ entangled state for 2-7 qubits; investigating a basic building block of general quantum computing. This is working toward the goal of running this in hardware on a superconducting flux qubit quantum annealing machine housed at the Quantum Artificial Intelligence Laboratory (QuAIL) at NASA’s Advanced Supercomputing facility.

Hans Briegel
University of Innsbruck

Projective simulation for learning and agency

I will first present the model of projective simulation (PS) for a learning agent, whose interaction with the environment is governed by a simulation-based projection. The PS agent uses random walks in its episodic and compositional memory (ECM) to project itself into future situations before taking real action. The PS model can solve basic tasks in reinforcement learning but it also allows for the implementation of advanced concepts such as generalization and meta-learning. Notably, projective simulation can be quantized, allowing for a quantum mechanical speed-up in the agent's deliberation process. I will then discuss recent applications of the PS model in robotics and the question to what extent learning agents can help us in finding new quantum experiments.

Vedran Dunjko
University of Innsbruck, Institute of Theoretical Physics

Quantum-enhanced learning

The field of artificial intelligence (AI) has lately had remarkable successes, specifically through progress in the field of classical machine learning. The emerging field of quantum machine learning has thus an even higher, and more profound potential to revolutionize the field of AI, much like quantum information processing has influenced many related but classical areas of research.

In this work we propose an approach for the systematic treatment of machine learning, from the perspective of quantum information. It is based on the agent-environment paradigm which we extend to a full quantum setting. Our methodology is general and covers all three main branches of machine learning: supervised, unsupervised and reinforcement learning.

While quantum improvements in supervised and unsupervised learning have been reported in other works, reinforcement learning has received much less attention. Within our approach, we tackle the problem of quantum enhancements in reinforcement learning as well, and propose a systematic scheme for providing improvements. Our basic idea relies on contrasting task environments to oracles, commonly used in classical and quantum information processing. We provide a series of examples, using our schematic approach, where quadratic improvements in learning efficiency, and an exponential improvements in performance over limited time periods, can be obtained for a broad class of learning problems.

<Invited Talk>

Joseph Fitzsimons
CQT, National University of Singapore

Statistical models and quantum information

In recent years, machine learning has emerged as an important application area for large scale quantum computers. This talk aims to explore some connections between the two fields. I will begin by introducing Gaussian processes, a statistical model commonly used in machine learning applications. I will discuss how quantum algorithms for Gaussian process regression can be built from existing quantum algorithms for matrix inversion and phase estimation. I will also discuss how techniques stemming from quantum information theory can be used to speed up purely classical algorithms in this area. Gaussian processes represent an entirely classical statistical model, whereas quantum systems allow for a much richer variety of correlations. I will conclude my talk by introducing a statistical model for quantum processes which extends the density matrix framework to capture quantum correlations throughout both space and time. This model shows a striking difference between the statistics of classical and quantum systems, as the statistics for measurements made at different points in a quantum process suffice to identify causal relationships, a task not possible with purely classical correlations.

Christina Giarmatzi
The University of Queensland

Finding a causal model for a set of classical variables is a well-established task but what about quantum variables? Even the notion of a quantum causal model is ambiguous. Here, we make use of a framework [Oreshkov, Costa, Bruckner, Nat. Comm. (2012)] that allows us to define causal models within quantum mechanics. In this framework, a causal model is represented by a Directed Acyclic Graph (DAG) in which the nodes and arrows represent quantum operations and quantum channels respectively. Making use of the process matrix formalism [Costa, Shrapnel, ArXiv (2015)] it was found that a single DAG imposes linear constraints on the process matrix, given that the process is Markovian. Since the process matrix can be calculated from measured probabilities when the parties perform an informationally complete set of operations, this provides a way that connects observational data to a DAG in the quantum case. However, reconstructing the DAG from a process matrix is an intricate and lengthy procedure, requiring a number of tests on the process matrix: each element of the DAG (an edge between two parties, which party is first or last etc) corresponds to a particular constraint on the matrix.

We present an algorithm that automates this set of actions. It takes as inputs: the number of parties, the dimensions of all input and output systems, if and how the output systems factorize into subsystems that are sent to different parties, and finally the process matrix. The algorithm outputs: a DAG describing all the causal relations between the parties, and which output subsystem (if any) of a party belongs to each causal arrow. Finally, the algorithm detects whether the process is Markovian or not, namely, whether all relevant common causes are included in the process, or else some causal relations are mediated through some external 'memory'. Our algorithm provides a first step towards more general methods for causal discovery of quantum variables.

Natalia Kaputkina
National University of Science and Technology MISIS (Russia)

Tuning quantum dot parameters with magnetic field.

Quantum dot (QD) are promising candidates for the hardware of future quantum computational devices, including quantum computers, quantum memory, quantum cellular automata, and quantum neural networks. The advantages of quantum dots are elaborated technology of QD production. In this work we consider QD arrays as prospective elements for quantum neural networks and discuss the means of controlling their properties with external electric and magnetic fields. Using QD arrays with typical dot size from few nanometers to tens nanometers and apply magnetic fields of tenths Tl order we can significantly change effective steepness potential of QDs and therefore the energy spectra, electron localization, tunneling coefficient between the dots. Influence of electrical field become considerable in fields about several V/cm. The increase of the external electric and the external magnetic field contributes to decrease of electron tunneling probability.

[1] Ya.M.Blanter, N.E.Kaputkina, and Yu.E.Loikov, Two-Electron Quantum Dots in Magnetic Field, Phys.Scripta v.54, pp.539-541, 1996.

[2] N.E.Kaputkina, Quantum Dots and Wells in External Electromagnetic Field, Chap.1 in Contemporary Research in Quantum Systems, ed. by Zoheir Ezziane, Nova Science Publishers, Incorporated, 2014.,363 pages.

[3] P.Yu.Korotaev, N.E.Kaputkina, Yu.E.Loikov, Yu.Kh.Vekilov, Electronic excitations and transport in aperiodic sequences of quantum dots in external electric and magnetic fields, JETP, τVol. 113. pp.692-697, 2011

[4] M. V. Altaisky, N. N. Zolnikova, N. E. Kaputkina, V. A. Krylov, Yu. E. Loikov, N. S. Dattani. Towards a feasible implementation of quantum neural networks using quantum dots. Appl. Phys. Lett. Vol. 108, p.103108, 2016

Zhaokai Li
University of Science and Technology of China

Experimental Realization of a Quantum Support Vector Machine

The fundamental principle of artificial intelligence is the ability of machines to learn from previous experience and do future work accordingly. In the age of big data, classical learning machines often require huge computational resources in many practical cases. Quantum machine learning algorithms, on the other hand, could be exponentially faster than their classical counterparts by utilizing quantum parallelism. Here, we demonstrate a quantum machine learning algorithm to implement handwriting recognition on a four qubit NMR test bench. The quantum machine learns standard character fonts and then recognizes handwritten characters from a set with two candidates. Because of the wide spread importance of artificial intelligence and its tremendous consumption of computational resources, quantum speedup would be extremely attractive against the challenges of big data.

Ref: (PhysRevLett.114.140504)

<Invited Talk>

Seth Lloyd
MIT

Quantum machine learning

This talk reviews algorithms for quantum machine learning, including cluster finding, support vector machines, principal component analysis, financial, and topological algorithms. The ability of quantum computers to perform Fourier transforms, invert matrices, and find eigenvectors and eigenvalues exponentially faster than their classical counterparts translates into highly efficient quantum machine learning algorithms.

Alex Monras
Universitat Autònoma de Barcelona

Structural risk minimization lies at the core of computational learning theory. Nevertheless, only few and sparse attempts have been made to consistently generalize it to the quantum scenario. We prove a de Finetti theorem for nonsignaling symmetric multipartite quantum channels, and use it to address the problem of structural risk minimization for quantum inductive learning protocols. In particular, we show that for any quantum learning task which addresses an arbitrary training set and a large number of problem instances, one can --up to an asymptotically vanishing extra risk-- replace a generic learning protocol by the sequence consisting on a) performing a measurement on the training set, and b) conditionally on the outcome, apply a given CPTP map on each one of the problem instances. This allows us to address fundamental questions in structural risk minimization and meaningfully enquire about model complexity and sample complexity in the quantum scenario. We show that under mild assumptions on the measurement of the training set, large deviation bounds hold for the expected risk, thus suggesting natural quantum generalizations of the Rademacher complexity and the VC dimension of a family of quantum classifiers.

<Invited Talk>

**Alejandro Perdomo-Ortiz
NASA Ames Research Center, USA**

TBA

<Invited Talk>

Patrick Reberstrost
MIT

Quantum matrix computations for machine learning and signal processing

<Invited Talk>

Mikel Sanz

University of the Basque Country

Quantum memristors and Neuromorphic quantum computing

Technology based on memristors, resistors with memory whose resistance depends on the history of the crossing charges, has lately enhanced the classical paradigm of computation with neuromorphic architectures. 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. Indeed, we provide numerical simulations showing that memory effects actually persist in the quantum regime. Our quantization method, specifically designed for superconducting circuits, may be extended to other quantum platforms, allowing for memristor-type constructions in different quantum technologies. The proposed quantum memristor is then, 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.

Maria Schuld
University of KwaZulu-Natal, South Africa

Classification on a quantum computer: Linear regression and ensemble methods

Quantum machine learning algorithms usually translate a machine learning methods into an algorithm that can exploit the advantages of quantum information processing. One approach is to tackle methods that rely on matrix inversion with the quantum linear system of equations routine. We give such a quantum algorithm based on unregularised linear regression. Opposed to closely related work from Wiebe, Braun and Lloyd [PRL 109 (2012)] our scheme focuses on a classification task and uses a different combination of core routines that allows us to process non-sparse inputs, and significantly improves the dependence on the condition number. The second part of the talk presents an idea that transcends the reproduction of classical results. Instead of considering a single trained classifier, practitioners often use ensembles of models to make predictions more robust and accurate. Under certain conditions, having infinite ensembles can lead to good results. We introduce a quantum sampling scheme that uses the parallelism inherent to a quantum computer in order to sample from 'exponentially large' ensembles that are not explicitly trained.

James Steck
Wichita State University

Machine Learning For Quantum Computation in Quantum Lattice Gas

J. Yepez, Vahala and Vahala and Soe show that a Bose-Einstein condensate, Type-II quantum lattice gas exhibits behavior that can be used for quantum computing. These physical systems are shown to behave as individual quantum bits, $\psi = (\text{up}, \text{down})$, in a 3D lattice according to the Schrödinger wave equation with $m=1/2$

$$i \hbar \dot{\psi} = -K \sigma \Delta^2 \psi + [g|\psi|^2 - a]\psi$$

On a domain Ω . The first term on the right hand side models spatial tunneling between neighboring qbits, controlled by the spatially varying parameter $K(x)$. The second term contains a nonlinear interaction controlled by the spatially varying parameters $g(x)$ and $a(x)$ (x is the spatial variable). $\Phi = (1,1) \cdot \psi$ We are using this as a quantum computer, programmed with quantum machine learning. Inputs are applied to defined regions in Ω , i.e. two small regions for a 2 input quantum gate. An output is measured on a separate region of Ω and the parameters K , g and a , are trained, rather than programmed, to give the correct computation. The probability amplitude ψ propagates through Ω as a quantum wave. J. Yepez et. al. have shown that the lattice can exhibit vortex and anti-vortex filaments and soliton type behavior. We investigate whether training produces values of K , a and g that create filaments and wave guides that serve as virtual synaptic pathways. In other words, training might create or “grow” pathways in the lattice that produce the desired computation. Preliminary simulations for a 2 input one output CNOT gate on a very coarse 2D mesh (40x50) verify our simulation and training algorithm. Results for a refined mesh applied to classical gates XOR, XNOR, CNOT, quantum gates, Toffoli, Fredkin, etc. will be ready for presentation at the conference.

Ryan Sweke
University of KwaZulu-Natal, South Africa

Tensor Network Methods, Neural Networks and Deep Learning

The last two decades has witnessed the emergence of tensor network methods as highly effective tools for both the analytical and numerical study of many-body quantum systems. In particular, tensor network representations of many-body states can be thought of as entanglement representations, and as a result can often be used to reduce the number of parameters necessary for the description of such states when specific correlation structures, such as area laws, are present. In a completely different direction, similar tensor decomposition approaches to those used in the many-body physics context have recently found a variety of applications in the fields of neural networks and machine learning. In this talk I will briefly survey the use of tensor decompositions in both the many-body physics as well as the machine learning context, with the goal of highlighting the potential for insights and techniques from one direction to influence the other.

Ramon Munoz Tapia
Universitat Autònoma de Barcelona

Sudden changes are ubiquitous in nature. In many physical settings there is a point when things start to be different from what they used to be. This may be due to a permanent alteration occurred at some previous time. Usually the time interval when the alteration happened is known or at least can be estimated. For many reasons we may be interested in knowing this specific time as this information has many significant practical consequences. In mathematical statistics this problem, the detection of sudden changes in the characteristics of the processes observed, is known as the change point problem. The change point can also be viewed as a border problem, where one wants to identify the line that separates different configurations and hence is also a central task in machine learning.

We propose a primitive quantum task that encapsulates this situation in its bare bones. A source assumed to prepare a sequence of identical states, starts to prepare a different states after some point. The length of the string of states where the change has occurred is given and the alteration is assumed to happen with equal probability at any point. We find the analytical expression of the optimal success probability of correct identification of the change point for large string lengths, which requires collective measurements on the whole string of states. We also analyse protocols that measure systems individually and provide an online answer at any stage of the measurement process and hence do not require the use of a quantum memory. We show that these underperform quite substantially the optimal collective procedure.

Peter Wittek
ICFO-The Institute of Photonic Sciences

Various flavours of causal networks -- Bayesian networks, Markov networks, conditional random fields, hidden Markov models and Kalman filters -- compactly represent the correlation or causal structure of a process under study by capturing the sparsity characteristics of the joint probability distribution. The graph thus encompasses the qualitative properties of the distribution. Unless this structure has a well-defined, simple form, probabilistic inference in causal networks is a challenging task. Markov networks, under mild conditions, correspond to Gibbs distributions, which allows inference by Gibbs sampling, which is normally performed using the Markov chain Monte Carlo (MCMC) method.

While it is straightforward to replace the classical MCMC method with a form of quantum Gibbs sampling, due to the arbitrary connectivity of the Markov network, the general problem remains a hard task. If we turn to Markov logic networks, however, the structure becomes more regular.

Markov logic networks (MLNs) reconcile the opposing schools in machine learning and artificial intelligence: one is causal networks, which account for uncertainty extremely well, and the other one is first-order logic, which comes from the symbolist tradition of artificial intelligence and relies on inverse deduction to perform inference. We can think of an MLN as a template to generate Markov networks, and the connectivity structure of the graph will depend on the maximum size of the formulas of a knowledge base defined in first-order logic.

This regular structure, opens up new possibilities to perform quantum Gibbs sampling either by using protocols with certified runtime and precision or by using quantum annealing. This offers quantum speed-ups and at the same time has great appeal to implementations.