Quantum Techniques in Machine Learning 2022
7 – 12 November, Naples, Italy

Program and Abstracts

Università degli Studi di Napoli Federico II

Università degli Studi di Napoli Federico II - Dipartimento di Fisica "Ettore Pancini"

QUNOVA COMPUTING

Università di Verona - Dipartimento di Informatica

data cybernetics

QUANTUMNET

QTML 2022 - Conference on Quantum Techniques in Machine Learning

	Monday 7 November	Tuesday 8 November	Wednesday 9 November	Thursday 10 November	Friday 11 November	Saturday 12 November
9:00 - 9:30		Conference Opening	Invited Speaker **Christa Zoufal**	Invited Speaker **Anna Sanpera**	Invited Speaker **Shaohan Hu**	***Workshop on Quantum Software***
9:30 - 10:20		Keynote Speaker ***Seth Lloyd***	Session *Variational quantum algorithms - Part I*	Session *Quantum Neural Networks and other topics*	Session *Quantum Learning - Part I*	
10:20 - 10:30					Sponsor Talk **QUANTUMNET**	
10:30 - 11:00		Coffee break	Coffee break	Coffee break	Coffee break	
11:00 - 12:00		Peter Wittek Session	Session *Variational quantum algorithms - Part II*	Session *Machine learning and Quantum systems - Part I*	Session *Quantum Learning - Part II*	
12:00 - 12:30		Keynote Speaker ***Gitta Kutyniok***	Keynote Speaker* ***Jacob Biamonte***	Invited Speaker **Marcello Dalmonte**	Invited Speaker **Shahnawaz Ahmed**	
12:30 - 13:00				Invited Speaker **Cristiano Ciuti**	Invited Speaker **Mario Krenn**	
13:00 - 14:00						
14:00 - 14:30		Tutorial speaker **Alex B. Grilo**	Invited Speaker **Florian Marquardt**	Industrial Track	Session *Quantum Learning - Part III*	
14:30 - 15:00			Session *Quantum Optimization*			
15:00 - 15:50	Conference Registration				Session *Computational Learning Theory*	
15:50 - 16:00			Sponsor talk **QUNOVA**			
16:00 - 16:30	Poster Session I	Coffee break	Coffee break	Coffee break		
16:30 - 17:00		Poster session II	Invited Speaker* **Patrick Coles**	Session *Machine learning and Quantum systems - Part II*	Coffee break	
17:00 - 17:30	Coffee break		Session *Quantum Neural Networks*		Invited Speaker* **Sabre Kais**	
17:30 - 18:30	Poster Session I			Keynote Speaker* ***Sankar Das Sarma***	Keynote Speaker* ***Hartmut Neven***	
20:00 - 21:30	Welcome reception		Conference Banquet			
21:30 - 22:30						

Workshop on Quantum Software

	Saturday 12 November
9:00 - 9:10	Opening
9:10 - 10:00	Keynote Speaker ***Mingsheng Ying***
10:00 - 10:20	Invited Speaker **Urs Schreiber**
10:20 - 10:40	Coffee break
10:40 - 11:00	Invited Speaker **Johanna Barzen**
11:00 - 11:20	Invited Speaker **Koen Bertels**
11:20 - 11:40	Invited Speaker **Sebastian Feld**
11:40 - 12:00	Invited speaker **Medina Bandic**
12:00 - 12:20	Invited Speaker **Kevin Mato**
12:20 - 13:10	Lunch break
13:10 - 13:30	Ilmo Salmenperä
13:30 - 13:50	Nicola Assolini
13:50 - 14:00	Closing

Contents

Presentation

Steering committee

General Chairs

Program Chairs

Abstracts of talks

Tutorial Talk: Alex B. Grilo

CNRS/Sorbonne Université, France

Introduction to quantum (statistical) learning theory

Given the key position of Machine Learning in everyday life, it is crucial to (formally) understand what can be efficiently learned or not. Learning theory is thus the field that studies learning problems from a complexity theory point of view, where we define formal models of learning and prove (in)feasibility results. With the advent of quantum computing, quantum speedups on learning tasks is a sought-after application. As in the classical setting, we are also interested in having a clear picture of what can be efficiently learned or not, especially when quantum advantage in such a setting is achieved. This is the goal of quantum learning theory. In this tutorial, I will present the field of learning theory, describing the models and main results. In particular, I will focus on quantum statistical learning theory, a model where data is accessed by its statistics. No previous knowledge of learning theory is expected.

Invited Talk: Christa Zoufal

IBM Zurich, Switzerland

Variational Quantum Imaginary Time Evolution: A discussion on error bounds, efficient approximations, and application examples

Variational quantum imaginary time evolution (VarQITE) allows us to simulate imaginary time dynamics of quantum systems with near-term compatible quantum circuits. These types of dynamics can be used to tackle a variety of important tasks such as searching for ground states, preparing Gibbs states, and optimizing black box binary optimization functions. In this talk, I will discuss several properties of this method and present examples executed with numerical simulations as well as actual quantum hardware. While variational approaches have the advantage of being compatible with short-depth quantum circuits, the underlying approximation errors are typically difficult to understand. I will explain how to find and efficiently evaluate a posteriori bounds for the approximation error of this variational technique. Furthermore, knowing the computational cost of a method as well as the related scaling behavior is important to understand what resources are required for the execution of an algorithm. This talk will also present an efficient approximation of VarQITE that reduces the cost significantly. Finally, I will present examples of VarQITE with a special focus on a real-world optimization problem, i.e. feature selection.

Session: Variational quantum algorithms - Part I

Exploiting symmetry in variational quantum machine learning

Johannes Jakob Meyer, Marian Mularski, Elies Gil-Fuster, Antonio Anna Mele, Francesco Arzani, Alissa Wilms and Jens Eisert

Variational quantum machine learning is an extensively studied application of near-term quantum computers. The success of variational quantum learning models crucially depends on finding a suitable parametrization of the model that encodes an inductive bias relevant to the learning task. However, precious little is known about guiding principles for the construction of suitable parametrizations. In this work, we holistically explore when and how symmetries of the learning problem can be exploited to construct quantum learning models with outcomes invariant under the symmetry of the learning task. Building on tools from representation theory, we show how a standard gateset can be transformed into an equivariant gateset that respects the symmetries of the problem at hand through a process of gate symmetrization. We benchmark the proposed methods on two toy problems that feature a non-trivial symmetry and observe a substantial increase in generalization performance. As our tools can also be applied in a straightforward way to other variational problems with symmetric structure, we show how equivariant gatesets can be used in variational quantum eigensolvers.

Avoiding barren plateaus via transferability of smooth solutions in Hamiltonian Variational Ansatz

Antonio Anna Mele, Glen Bigan Mbeng, Giuseppe Ernesto Santoro, Mario Collura and Pietro Torta

A large ongoing research effort focuses on Variational Quantum Algorithms (VQAs), representing leading candidates to achieve computational speed-ups on current quantum devices. The scalability of VQAs to a large number of qubits, beyond the simulation capabilities of classical computers, is still debated. Two major hurdles are the proliferation of low-quality variational local minima, and the exponential vanishing of gradients in the cost function landscape, a phenomenon referred to as barren plateaus. Here we show that by employing iterative search schemes one can effectively prepare the ground state of paradigmatic quantum many-body models, circumventing also the barren plateau phenomenon. This is accomplished by leveraging the transferability to larger system sizes of iterative solutions, displaying an intrinsic smoothness of the variational parameters, a result that does not extend to other solutions found via random-start local optimization. Our scheme could be directly tested on near-term quantum devices, running a refinement optimization in a favorable local landscape with non-vanishing gradients.

Quantum variational learning for quantum error-correcting codes

Chenfeng Cao, Chao Zhang, Zipeng Wu, Markus Grassl and Bei Zeng

Quantum error-correcting codes (QECCs) are believed to be a necessity for large-scale fault-tolerant quantum computation. In the past two decades, various methods of QECC constructions have been developed, leading to many good families of codes. However, the majority of these codes are not suitable for near-term quantum devices. Here we present VarQEC, a noise-resilient variational quantum algorithm to search for quantum codes with a hardware-efficient encoding circuit. The cost functions are inspired by the most general and fundamental require- ments of a QECC, the Knill-Laflamme conditions. Given the target noise channel (or the target code parameters) and the hardware connectivity graph, we optimize a shallow variational quantum circuit to prepare the basis states of an eligible code. In principle, VarQEC can find quantum codes for any error model, whether additive or non-additive, degenerate or non-degenerate, pure or impure. We have verified its effectiveness by (re)discovering some symmetric and asymmetric codes, e.g., $((n, 2^{n-6}, 3))_2$ for n from 7 to 14. We also found new $((6, 2, 3))_2$ and $((7, 2, 3))_2$ codes that are not equivalent to any stabilizer code, and ex- tensive numerical evidence with VarQEC suggests that a $((7, 3, 3))_2$ code does not exist. Furthermore, we found many new channel-adaptive codes for error models involving nearest-neighbor correlated errors. Our work sheds new light on the understanding of QECC in general, which may also help to enhance near-term device performance with channel-adaptive error-correcting codes. The full article can be found here: https://arxiv.org/pdf/2204.03560.pdf

Session: Variational quantum algorithms - Part II

Select your metric carefully: achieving provably efficient variational modeling

Faris Sbahi, Antonio Martinez, Sahil Patel, Dmitri Saberi, Jae Hyeon Yoo, Geoffrey Roeder and Guillaume Verdon

The dual tasks of quantum Hamiltonian learning and quantum Gibbs sampling are relevant to many important problems in physics and chemistry. In the low temperature regime, algorithms for these tasks often suffer from intractabilities, for example from poor sample- or time-complexity. With the aim of addressing such intractabilities, we introduce a generalization of quantum natural gradient descent to parameterized mixed states, as well as provide a robust first-order approximating algorithm, Quantum-Probabilistic Mirror Descent. We prove data sample efficiency for the dual tasks using tools from information geometry and quantum metrology, thus generalizing the seminal result of classical Fisher efficiency to a variational quantum algorithm for the first time.

Policy Gradients using Variational Quantum Circuits

André Sequeira, Luis Paulo Santos and Luis Soares Barbosa

Variational Quantum Circuits are being used as versatile Quantum Machine Learning models. Some empirical results exhibit an advantage in supervised and generative learning tasks. However, when applied to Reinforcement Learning, less is known. In this work, we considered a Variational Quantum Circuit composed of a low-depth hardware-efficient ansatz as the parameterized policy of a Reinforcement Learning agent. We show that an σ-approximation of the gradient can be obtained using a logarithmic number of samples concerning the total number of parameters and quantum circuit executions. We empirically verify that such quantum models behave similarly or even outperform typical classical neural networks used in benchmarking environments with a fraction of the parameters. Moreover, we study the Barren Plateau phenomenon in policy gradients using the Fisher Information Matrix trace. Furthermore, we applied the variational quantum model to a quantum control problem and showed its feasibility in the quantum-quantum domain compared to classical Reinforcement Learning.

Policy Gradient for Training Variational NISQ Algorithms

David Herrera Marti

We propose a method for training variational quantum circuits, based on techniques from policy gradient reinforcement learning. The choice of a stochastic policy allows us to rephrase the optimization problem in terms of probability distributions, rather than variational gates. In this framework, finding the optimal configuration is done by optimizing over distribution parameters, rather than over free angles. We show numerically that this approach can be more competitive than gradient-free methods, for comparable amounts of resources (i.e. quantum circuit runs), at the onset of a barren plateau.

An analogue quantum variational embedding classifier

Adrian Lupascu, Rui Yang, Samuel Bosch, Bobak Kiani and Seth Lloyd

Quantum machine learning has the potential to provide powerful algorithms for artificial intelligence. The pursuit of quantum advantage in quantum machine learning is an ongoing area of research. For current noisy, intermediate-scale quantum (NISQ) computers, various quantum-classical hybrid algorithms have been proposed. One such hybrid algorithm is a gate- based variational embedding classifier, which is composed of a classical neural-network and a gate-based quantum circuit. We propose a quantum variational embedding classifier based on an analogue quantum computer, with a particular focus on a quantum annealer implementation. In our classifier, the nonlinear mapping of the classical data to a high-dimensional density matrix is realized in an analogue manner. The nonlinearity needed for a nonlinear classification problem is purely provided by the analogue quantum computer. We demonstrate the effectiveness of our algorithm for performing binary and multilabel classification on linearly inseparable datasets, including concentric rings and MNIST digits. We also characterized the performance of our algorithm in terms of the error scaling vs the number of qubits and the dependence on running time. Our algorithm presents the possibility to use current quantum annealers for solving practical machine-learning problems and it could also be useful to explore quantum advantage in quantum machine learning.

Characterization of variational quantum algorithms using free fermions

Gabriel Matos, Chris N. Self, Zlatko Papic, Konstantinos Meichanetzidis and Henrik Dreyer

We study variational quantum algorithms from the perspective of free fermions. Using Lie theoretical techniques, we characterize the space of states that the Quantum Approximate Optimization Algorithm (QAOA) is able to prepare at arbitrary circuit depth on a one-dimensional lattice with and without decoupled angles. We show that this is the set of all fermionic Gaussian states respecting the symmetries of the circuit, and we numerically determine the minimum depth necessary to prepare any such state. In several cases we find that these protocols can be made into exact parameterizations of this space, so that global minima exist and are unique. We proceed to study the interplay between the symmetries of the circuit and the locality of the target state and find that an absence of symmetries makes nonlocal states easier to prepare. An efficient classical simulation of Gaussian states is leveraged in order to expose behavior not evident at smaller size. We exploit this in order to study the behavior of the overparameterized regime of optimization with increasing circuit depth, and to show that it saturates quadratically with the system size.

Qubit seriation: Undoing data shuffling using a spectral ordering approach

Atithi Acharya, Manuel Rudolph, Jing Chen and Alejandro Perdomo Ortiz

With the advent of quantum and quantum-inspired machine learning, designing the model structure for particular datasets has been shown to be crucial for success. Tensor network-based learning models are prime candidates to benefit from careful design considerations since these models are biased towards data with local correlations matching the structure of the model's topology. In our work, we aim to find optimal permutations of sites in the data by utilizing methods from spectral graph theory. This keeps strongly correlated bits closer with respect to the model topology, and weakly correlated bits farther away. The prepossessing is done based on pairwise mutual information extracted from the data. Here, we demonstrate the power of such prepossessing,by substantially improving training in the performance of TN models, specifically Matrix Product States-based generative models. We also gain further insights into the structure of the original data by using spectral embedding, a dimensionality reduction technique also originating in spectral graph theory.

Keynote Talk: Jacob Biamonte

Skolkovo Institute of Science and Technology, Russian Federation

Subtle effects in quantum circuits as machine learning models

I will explain some recent findings by myself and in collaboration: the discovery of several limiting features of quantum approximate optimisation (reachability deficits and parameter training saturations) and the existence of abrupt trainability transitions in quantum circuit training. I will also explain some more forward looking findings, including the concentration of parameters in quantum approximate optimisation (showing a problem instance independence of optimised circuit parameters) and my proof that the variational model is, in theory, a universal model of quantum computation.

Invited Talk: Florian Marquardt

Max Planck Institute, Germany

Combining Gradient Ascent and Feedback Control

Optimal control algorithms are essential for improving modern quantum devices. While model-based gradient techniques like GRAPE (gradient-ascent pulse engineering) are powerful tools for efficiently finding control pulses, they are not applicable to feedback scenarios, where the control must depend on measurement results. Conversely, modern model-free reinforcement learning techniques can easily deal with feedback, but they are not very efficient, since they do not make use of our knowledge of the underlying physics model. In this talk, I will present our new approach (termed feedback-GRAPE) that enables us to combine model-based techniques with quantum feedback. I will give examples of several tasks that can be efficiently solved using that new approach.

Session: Quantum Optimization

GEO: Enhancing Combinatorial Optimization with Classical and Quantum Generative Models

Francisco Javier Fernandez Alcazar, Mohammad Ghazi Vakili, Can B. Kalayci and Alejandro Perdomo-Ortiz

We introduce a new framework that leverages machine learning models known as generative models to solve optimization problems. Our Generator-Enhanced Optimization (GEO) strategy is flexible to adopt any generative model, from quantum to quantum-inspired or classical, such as Generative Adversarial Networks, Variational Autoencoders, or Quantum Circuit Born Machines, to name a few. Here, we focus on a quantum-inspired version of GEO relying on tensor-network Born machines, and referred to hereafter as TN-GEO. We present two prominent strategies for using TN-GEO. The first uses data points previously evaluated by any quantum or classical optimizer, and we show how TN-GEO improves the performance of the classical solver as a standalone strategy in hard-to-solve instances. The second strategy uses TN-GEO as a standalone solver, i.e., when no previous observations are available. Here, we show its superior performance when the goal is to find the best minimum with the least number of cost function evaluations. This might be ideal in situations where the cost function evaluation can be very expensive. To illustrate our results, we run these benchmarks in the context of the portfolio optimization problem by constructing instances from the S&P 500 and several other financial stock indexes. We show that TN-GEO can propose unseen candidates with lower cost function values than the candidates seen by classical solvers. This is the first demonstration of the generalization capabilities of quantum-inspired generative models that provide real value in the context of an industrial application. We also comprehensively compare state-of-the-art algorithms in a generalized version of the portfolio optimization problem. The results show that TN-GEO is among the best compared to these state-of-the-art algorithms; a remarkable outcome given the solvers used in the comparison have been fine-tuned for decades in this real-world industrial application. We see this as an important step toward a practical advantage with quantum-inspired models and, subsequently, with quantum generative models.

Benchmarking Quantum Heuristics and Ising Machines: Scoring Parameter Setting Strategies

Davide Venturelli, David Bernal Neira, Filip Wudarski and Eleanor Rieffel

We discuss guidelines for the performance evaluation of parameterized stochastic solvers of optimization problems, with particular attention to systems that employ quantum or unconventional processors, such as QAOA, Quantum Annealing, or Coherent Ising Machines. We illustrate through an example a benchmarking procedure grounded in the statistical analysis of the expectation of a given performance metric outside a test environment, taking into account all resources related to real-world deployment. By properly discussing these aspects, it is possible to identify evaluation principles of parameter tuning strategies and more complex use of the heterogeneous components of the solver. In particular we examine in detail and provide the software to compute the scoring of an illustrative baseline family of parameter setting strategies that feature an exploration-exploitation tradeoff, on a parallel tempering solver tackling hard spin-glass instances, discovering that the best parameter setting strategy among those studied employs Bayesian estimation to search through the parameter space.

Qubit-efficient encoding schemes for binary optimisation problems

Benjamin Tan, Marc-Antoine Lemonde, Supanut Thanasilp, Jirawat Tangpanitanon and Dimitris Angelakis

We propose and analyse a set of variational quantum algorithms for solving quadratic unconstrained binary optimization problems which can be implemented on $O(log(n))$ number of qubits. The underlying encoding scheme allows for a systematic increase in correlations among the classical variables captured by a variational quantum state by progressively increasing the number of qubits involved. We first examine the simplest limit where all correlations are neglected, i.e. when the quantum state can only describe statistically independent classical variables. We apply this minimal encoding to find approximate solutions of a general problem instance comprised of 64 classical variables using 7 qubits. Next, we show how two-body correlations between the classical variables can be incorporated in the variational quantum state and how it can improve the quality of the approximate solutions. We give an example by solving a 42-node Max-Cut problem using only 8 qubits where we exploit the specific topology of the problem. We analyze whether these cases can be optimized efficiently given the limited resources available in state-of-the-art quantum platforms. Lastly, we present the general framework for extending the expressibility of the probability distribution to any multi-body correlations.

Quantum Annealing for Neural Network optimization problems: a new approach via Tensor Network simulations

Guglielmo Lami, Pietro Torta, Giuseppe Santoro and Mario Collura

Quantum Annealing (QA) is one of the most promising frameworks for quantum optimization. Here, we focus on the problem of minimizing complex classical cost functions associated with prototypical discrete neural networks, specifically the paradigmatic Hopfield model and binary perceptron. We show that the adiabatic time-evolution of QA can be efficiently represented as a suitable Tensor Network. This representation allows for simple classical simulations, well-beyond small sizes amenable to exact diagonalization techniques. We show that the optimized state, expressed as a Matrix Product State (MPS), can be recast into a Quantum Circuit, whose depth scales only linearly with the system size and quadratically with the MPS bond dimension. This may represent a valuable starting point allowing for further circuit optimization on near- term quantum devices.

Genetic Algorithms for Constructing Quantum Embeddings

Massimiliano Incudini, Francesco Martini and Alessandra Di Pierro

Quantum kernels require an appropriate choice of a quantum feature map, or quantum embedding, to enhance the performance of a distance-based classifier for supervised learning tasks. However, we rarely have enough knowledge of the data to allow us the construction of a satisfying structure for our embedding. We propose the use of genetic algorithms to to search for the best structure that maximises the Kernel-Target Alignment (a proxy metric for the accuracy of the model). We show experimentally that our approach outperforms the trainable quantum embedding approach based on gradient-descent optimization with respect to an artificial dataset.

Quantum Neural Architecture Search with Quantum Circuits Metric and Bayesian Optimization

Trong Duong, Sang Truong, Minh Nguyen, Bao Bach, Ju-Young Ryu and June-Koo Rhee

Quantum neural networks are promising for a wide range of applications in the Noisy Intermediate-Scale Quantum era. As such, there is an increasing demand for automatic quantum neural architecture search. We tackle this challenge by designing a quantum circuits metric for Bayesian optimization with Gaussian process. To this goal, we propose a new quantum gates distance that characterizes the gates' action over every quantum state and provide a theoretical perspective on its geometrical properties. Our approach significantly outperforms the benchmark on three empirical quantum machine learning problems including training a quantum generative adversarial network, solving combinatorial optimization in the MaxCut problem, and simulating quantum Fourier transform. Our method can be extended to characterize behaviors of various quantum machine learning models.

Invited Talk: Patrick Coles

The University of New Mexico, New Mexico

Trainability and generalization of quantum machine learning models

Quantum machine learning (QML) is a leading proposal for near-term quantum advantage. However, recent progress in understanding the training landscapes for QML models paints a concerning picture. Exponentially vanishing gradients-barren plateaus-occur for circuits that are deep or noisy or generate much entanglement. On the flip side, some QML architectures are immune to barren plateaus. The dynamical Lie algebra has been connected to both barren plateaus and overparameterization, suggesting that we could use algebraic properties to engineer favorable training landscapes. Training is only half of the story for QML, as good generalization to testing data is needed as well. We have found surprising good generalization properties for QML models in two ways: (1) In-distribution generalization is guaranteed when the training data size is roughly equal to the number of model parameters, and (2) Out-of-distribution generalization is guaranteed for locally scrambled ensembles, allowing for product-state to entangled-state generalization. In this talk, I will attempt to overview our current understanding of both the trainability and generalization of QML models.

Session: Quantum Neural Networks

Benign overfitting in shallow quantum neural networks

Evan Peters and Maria Schuld

The widespread success of deep neural networks has revealed a surprise in classical machine learning: very complex models often generalize well while simultaneously overfitting training data. This phenomenon of benign overfitting has been studied for a variety of classical models with the goal of better understanding the mechanisms behind deep learning. Understanding this phenomenon in the context of quantum machine learning might similarly improve our understanding of the relationship between overfitting, overparameterization, and generalization. In this work, we provide a characterization of benign overfitting in quantum models. To do this, we derive the behavior of a classical interpolating Fourier features models for regression on noisy signals, and show how a class of quantum models exhibits analogous structure, thereby linking data-encoding and state preparation in quantum models to over-parameterization and overfitting. We intuitively explain these behaviors according to the ability of the quantum model to interpolate noisy data with locally "spiky" behavior and provide a concrete demonstration example of benign overfitting.

Ensembling Techniques for Quantum Neural Networks

Massimiliano Incudini, Michele Grossi, Andrea Ceschini, Antonio Mandarino, Massimo Panella, Sofia Vallecorsa, David Windridge and Alessandra Di Pierro

Quantum Neural Networks are, in principle, capable of solving supervised learning tasks in respect to both classical and quantum data with potential benefits in terms of generalization error. However, current noisy and resource-limited quantum hardware typically fails to solve real-world machine learning problems. We here propose a framework for building classical ensembles of Quantum Neural Networks in order to mitigate current hardware limitations. We find experimentally that bagging ensembles are indeed able both to reduce the number of qubits needed to solve a task and also to improve the generalization error on synthetic noisy linear data, with potential additional benefits in relation to the barren plateau problem.

Introducing Non-Linearity into Quantum Generative Models

Mykolas Sveistrys, Kaitlin Gili and Chris Balance

Current quantum generative models use circuits that are unitary and hence reversible, due to the linear nature of quantum mechanics. However, classical neural-network-based generative models involve highly non-linear and thus non-reversible dynamics. In this paper, we explore the effect of these dynamics in quantum generative modeling by introducing a model that adds non-linear activations via a neural network structure onto the standard Born Machine framework - the Quantum Neuron Born Machine (QNBM). To achieve this, we utilize a previously introduced Quantum Neuron subroutine, which is a repeat-until-success circuit with mid-circuit measurements and classical control. We compare our non-linear QNBM to the linear Quantum Circuit Born Machine (QCBM). We allocate similar time and memory resources to each model, such that the only major difference is the qubit overhead required by the QNBM. With gradient-based training, we show that while both models can easily learn a trivial uniform probability distribution, on a more challenging class of distributions, the QNBM achieves an almost 3x smaller error rate than a QCBM with a similar number of tunable parameters. We therefore provide evidence that suggests that non-linearity is a useful resource in quantum generative models, and we put forth the QNBM as a new model with good generative performance and potential for quantum advantage.

Entanglement entropy production in Quantum Neural Networks

Stefano Mangini, Marco Ballarin, Simone Montangero, Chiara Macchiavello and Riccardo Mengoni

Quantum Neural Networks (QNN) are considered a candidate for achieving quantum advantage in the Noisy Intermediate Scale Quantum computer (NISQ) era. Several QNN architectures have been proposed and successfully tested on benchmark datasets for machine learning. However, quantitative studies of the QNN generated entanglement have not been investigated in details, and only for up to few qubits. Tensor network methods allow to emulate quantum circuits with a large number of qubits in a wide variety of scenarios. Here, we employ matrix product states to characterize recently studied QNN architectures with up to fifty qubits showing that their entanglement, measured in terms of entanglement entropy between qubits, tends to that of Haar distributed random states as the depth of the QNN is increased. We show a universal behavior for the entanglement entropy production for any given QNN architecture, consequently we introduce a new measure to characterize the entanglement production in QNNs: the entangling speed. Finally, in agreement with known results in the literature, we argue that the most promising regime for quantum advantage with QNNs is defined by a trade-off between high entanglement and expressibility

Rapid training of quantum recurrent neural network

Michał Siemaszko, Thomas McDermott, Adam Buraczewski, Bertrand Le Saux and Magdalena Stobińska

Time series prediction is the crucial task for many human activities e.g. weather forecasts or predicting stock prices. One solution to this problem is to use Recurrent Neural Networks (RNNs). Although they can produce accurate predictions, their learning process is slow and complex. Here, we propose a Quantum Recurrent Neural Network (QRNN) to address these obstacles. The design of the network is based on the continuous-variable quantum computing paradigm. We demonstrate that the network is capable of learning time dependence of a few types of temporal data. Our numerical simulations show that the QRNN converges to optimal weights in fewer epochs than the classical network. Furthermore, for a small number of trainable parameters, it can achieve a loss lower than that of the latter.

QuanGCN: Noise-Adaptive Training for Robust Quantum Graph Convolutional Networks

Kaixiong Zhou, Zhenyu Zhang, Shengyuan Chen, Tianlong Chen, Xiao Huang, Zhangyang Wang and Xia Hu

Quantum neural networks(QNNs),an interdisciplinary field of quantum computing and machine learning, have attracted tremendous research interests due to the specific quantum advantages. Despite lots of efforts developed in computer vision domain, one has not fully explored QNNs for the real-world graph property classification and evaluated them in the quantum device. To bridge the gap, we propose quantum graph convolutional networks (QuanGCN), which learns the local message passing among nodes with the sequence of crossing-gate quantum operations. To mitigate the inherent noises from modern quantum devices, we apply sparse constraint to sparsify the nodes' connections and relieve the error rate of quantum gates, and use skip connection to augment the quantum outputs with original node features to improve robustness. The experimental results show that our QuanGCN is functionally comparable or even superior than the classical algorithms on several benchmark graph datasets. The comprehensive evaluations in both simulator and real quantum machines demonstrate the applicability of QuanGCN to the future graph analysis problem.

The complexity of quantum support vector machines

Gian Gentinetta, Arne Thomsen, David Sutter and Stefan Woerner

Quantum support vector machines employ quantum circuits to define the kernel function. It has been shown that this approach offers a provable exponential speedup compared to any known classical algorithm for certain data sets. The training of such models corresponds to solving a convex optimization problem either via its primal or dual formulation. Due to the probabilistic nature of quantum mechanics, the training algorithms are affected by statistical uncertainty, which has a major impact on their complexity. We show that the dual problem can be solved in $O(M^{4.67}/\epsilon^2)$ quantum circuit evaluations, where M denotes the size of the data set and ϵ the solution accuracy. We prove under an empirically motivated assumption that the kernelized primal problem can alternatively be solved in $O(min\{M^2/\epsilon^6, 1/\epsilon^{10}\})$ evaluations by employing a generalization of a known classical algorithm called Pegasos. Accompanying empirical results demonstrate these analytical complexities to be essentially tight. In addition, we investigate a variational approximation to quantum support vector machines and show that their heuristic training achieves considerably better scaling in our experiments.

Invited Talk: Anna Sanpera

Universitat Autonoma de Barcelona, Spain

Coherence theory in attractor quantum neural networks

One way to understand quantum neural networks is to adapt classical cases into the quantum regime. Attractor neural networks are able to retrieve different configurations after they are applied several times allowing to associate each initial state with the closest stable configuration of the network. The quantum case is obtained by studying which are the completely positive trace preserving (CPTP) maps that hold the larger number of stationary states. I will show that in this case, the When talking states This is done by can use a classical attractor neural network of We study the storage capacity of quantum neural networks (QNNs), described by the attractor associated to an arbitrary input state is the one minimizing their relative entropy. We will discuss why this networks outperform the classical ones.

Session: Quantum Neural Networks and other topics

Quantum Machine Learning for Portfolio Hedging Optimization

Jules Plaire, Paulin Jacquot and Thomas Deschatre

Financial portfolio hedging is a complex mathematical problem with no analytical solution in general, and for which classical machine learning methods have already been explored. We propose a hybrid classical-quantum optimization algorithm based on Parametrized Quantum Circuits (PQC) and inspired by reinforcement learning to optimize the hedging strategy of a call option in a single asset market. Based on simulations, we observe that PQCs converge faster and achieve similar performance than classical neural networks with a comparable number of parameters.

Evaluation of variational quantum states entanglement on a quantum computer by the mean value of spin

Khrystyna Gnatenko

The geometric measure of entanglement of variational quantum states is studied on the basis of its relation with the mean value of spin. We examine n-qubit quantum states prepared by a variational circuit with a layer formed by the rotational gates and two-qubit controlled phase gates. The variational circuit is a generalization of that used for preparing quantum Generative Adversarial Network states. The entanglement of a qubit with other qubits in the variational quantum states is determined by the angles of rotational gates that act on the qubit and qubits entangled with it by controlled phase gates and also their parameters. In the case of one layer variational circuit, the states can be associated with graphs with vertices representing qubits and edges corresponding to two-qubit gates. The geometric measure of entanglement of a qubit with other qubits in the quantum graph state depends on the properties of the vertex that represents it in the graph, namely it depends on the vertex degree. The dependence of the geometric measure of entanglement of variational quantum states on their parameters is quantified on IBM's quantum computer.

Efficient Quantum State Preparation Using Decision Diagrams

Fereshte Mozafari, Giovanni De Micheli and Yuxiang Yang

Loading classical data into quantum registers is one of the most important primitives of quantum computing. While the complexity of preparing a generic quantum state is exponential in the number of qubits, in many practical tasks the state to prepare has a certain structure that allows for faster preparation. In this paper, we consider quantum states that can be efficiently represented by (reduced) decision diagrams, a versatile data structure for the representation and analysis of Boolean functions. We design an algorithm that utilises the structure of decision diagrams to prepare their associated quantum states. Our algorithm has a circuit complexity that is linear in the number of paths in the decision diagram. Numerical experiments show that our algorithm reduces the circuit complexity by up to 31.85% compared to the state-of-the-art algorithm, when preparing generic n-qubit states with different degrees of non-zero amplitudes. Additionally, for states with sparse decision diagrams, including the initial state of the quantum Byzantine agreement protocol, our algorithm reduces the number of CNOTs by 86.61% - 99.9%.

Conditional Born machine for Monte Carlo events generation

Oriel Kiss, Michele Grossi, Enrique Kajomovitz and Sofia Vallecorsa

Generative modeling is a promising task for near-term quantum devices, which can use the stochastic nature of quantum measurements as random source. So called Born machines are purely quantum models and promise to generate probability distributions in a quantum way, inaccessible to classical computers. This paper presents an application of Born machines to Monte Carlo simulations and extends their reach to multivariate and conditional distributions. Models are run on (noisy) simulators and IBM Quantum superconducting quantum hardware. More specifically, Born machines are used to generate muonic force carriers (MFC) events resulting from scattering processes between muons and the detector material in high-energy-physics colliders experiments. MFCs are bosons appearing in beyond the standard model theoretical frameworks, which are candidates for dark matter. Empirical evidences suggest that Born machines can reproduce the underlying distribution of datasets coming from Monte Carlo simulations, and are competitive with classical machine learning-based generative models of similar complexity. The complete paper can be found in Ref. [1].

Monitored Quantum Reservoir Computing

Pere Mujal, Rodrigo Martínez-Peña, Gian Luca Giorgi, Miguel C. Soriano and Roberta Zambrini

Nowadays, classical machine learning algorithms are a necessary tool to process the increasing volume of available big data, including temporal sequences. Prominent time-series processing tasks are speech recognition or stock market and climate forecasting. In this context, a neuromorphic approach known as reservoir computing has become popular. The development of its quantum counterpart, quantum reservoir computing, has given rise to theoretical proposals that display superior performances in different platforms, including current noisy intermediate-scale quantum computers. However, a central issue towards the experimental implementation of online time-series processing is the destructive effect of quantum measurements. We tackle this problem by proposing and analyzing different realistic measurement protocols still achieving high performance, enabled by quantum coherence, and efficiency. Detrimental measurement effects have not been addressed before, assuming a reset of the system at each new input. We show how to improve the efficiency by harnessing the reservoir fading memory or also by considering a genuine online protocol based on weak measurements. Through analytical and numerical analysis, we compare the experimental time resources necessary to exploit the memory and the predictive capacity of quantum reservoir computing and establish strategies to efficiently reach the ideal performance. This analysis, beyond ideal situations and taking time (and energy) resources into account, is essential to stimulate experimental developments. Furthermore, we set a framework to be applied also in other monitored setups from temporal tomography to neuromorphic and recurrent neural networks implementations.

Quantum Fuzzy Control of Particle Beams

Giovanni Acampora, Michele Grossi, Michael Schenk and Roberto Schiattarella

Recently a quantum fuzzy inference engine that achieves an exponential advantage in computing fuzzy rules with respect to the classical counterpart has been proposed. The main goal of such a quantum engine is to pave the way for the implementation of fuzzy rule-based systems in complex environments where the number of fuzzy rules to compute is impractical for classical fuzzy systems. However, due to the technological limitation of current NISQ devices, the quantum engine has been tested only in noiseless simulations and in simplified environments. In this work, the quantum system is tested in controlling a simplified model of the T4 target station at the CERN SPS fixed target physics beam line and for the very first time, the quantum engine is computed on IBM quantum hardware. Therefore, this paper achieves a twofold goal: to show the suitability of this quantum fuzzy inference engine in controlling such environments and to provide proof of principle for its implementation on real quantum hardware.

Session: Machine learning and Quantum systems - Part I

Conceptual understanding enabled by efficient automated design of quantum optical setups

Mario Krenn, Jakob Kottmann, Nora Tischler and Alan Aspuru-Guzik

Artifcial intelligence is a powerful tool for science, but an important question is how to extract true scientific understanding. We present a method that enables new understanding, and demonstrate its application to quantum photonics. To achieve this, we make four main contributions: (i) We introduce a graph-based representation of photonic quantum experiments amenable to interpretation and algorithmic use. (ii) We develop an approach for the automated design of new quantum experiments, for which our benchmarking establishes a speed-up by orders of magnitude compared to previous algorithms. (iii) We solve several open questions in experimental quantum photonics; the solutions include blueprints for quantum states that are important for photonic quantum technology and for fundamental quantum experiments. (iv) The interpretable representation and massive speed-up allow us to obtain solutions from where scientists can extract new scientific concepts.

Neural network enhanced measurement efficiency for molecular groundstates

Dmitri Iouchtchenko, Jérôme F. Gonthier, Alejandro Perdomo-Ortiz and Roger G. Melko

It is believed that one of the first useful applications for a quantum computer will be the preparation of groundstates of molecular Hamiltonians. A crucial task involving state preparation and readout is obtaining physical observables of such states, which are typically estimated using projective measurements on the qubits. At present, measurement data is costly and time-consuming to obtain on any quantum computing architecture, which has significant consequences for the statistical errors of estimators. We adapt common neural network models (restricted Boltzmann machines and recurrent neural networks) to learn complex groundstate wavefunctions for several prototypical molecular qubit Hamiltonians from typical measurement data. By relating the accuracy ε of the reconstructed groundstate energy to the number of measurements, we find that using a neural network model provides a robust improvement over using single-copy measurement outcomes alone to reconstruct observables. This enhancement yields an asymptotic scaling near ε^{-1} for the model-based approaches, as opposed to ε^{-2} in the case of classical shadow tomography.

Out-of-distribution generalization for learning quantum dynamics and dynamical simulation

Matthias C. Caro, Hsin-Yuan Huang, Joe Gibbs, Nicholas Ezzell, Andrew Sornborger, Lukasz Cincio, Patrick Coles and Zoe Holmes

Generalization bounds are a critical tool to assess the training data requirements of Quantum Machine Learning (QML). In this work, we prove the first out-of-distribution generalization guarantees in QML, where we require a trained model to perform well even on testing data drawn from a distribution different from the training data distribution. Namely, we establish out-of-distribution generalization for the task of learning an unknown unitary using a quantum neural network and for a broad class of training and testing distributions. In particular, we show that one can learn the action of a unitary on entangled states using only product state training data. Since product states can be prepared using only single-qubit gates, this advances the near-term prospects of QML for learning quantum dynamics, and further opens up new methods for both the classical and quantum compilation of quantum circuits. Based on these insights, we propose a QML-based algorithm for simulating quantum dynamics on near-term quantum hardware and rigorously prove its resource-efficiency in terms of qubit and training data requirements. We also demonstrate the viability of this algorithm through numerical experiments, both in classical simulations and on quantum hardware. Finally, we embed this algorithm in a broader framework for using QML methods for quantum dynamical simulation on NISQ devices.

Invited Talk: Marcello Dalmonte

International Centre for Theoretical Physics, Italy

Data mining the output of quantum simulators - from critical behavior to algorithmic complexity

Recent experiments with quantum simulators and noisy intermediate-scale quantum devices have demonstrated unparalleled capabilities of probing many-body wave functions, via directly probing them at the single quantum level via projective measurements. However, very little is known about to interpret and analyse such huge datasets. In this talk, I will show how it is possible to provide such characterisation of many-body quantum hardware via a direct and assumption-free data mining. The core idea of this programme is the fact that the output of quantum simualtors and computers can be construed as a very high-dimensional manifold. Such manifold can be characterised via basic topological concepts, in particular, by their intrinsic dimension. Exploiting state of the art tools in non-parametric learning, I will discuss theoretical results for both classical and quantum many-body spin systems that illustrate how data structures undergo structural transitions whenever the underlying physical system does, and display universal (critical) behavior in both classical and quantum mechanical cases. I will conclude with remarks on the applicability of our theoretical framework to synthetic quantum systems (quantum simulators and quantum computers), and emphasize its potential to provide a direct, scalable measure of Kolmogorov complexity of output states.

Invited Talk: Cristiano Ciuti

Université Paris Cité, France

Exploiting machine learning for quantum dynamics and viceversa

This talk is about the recent theoretical research of our group on the reciprocal link between machine learning and the dynamics of quantum systems. First, we will review how artificial neural networks can be used as variational trial wavefunctions to simulate closed and open quantum systems. Then, we will show how the dynamics of quantum hardware can be exploited to create kernel machines to perform advanced tasks.

Session: Machine learning and Quantum systems - Part II

Flexible learning of quantum states with generative query neural networks

Yan Zhu, Ya-Dong Wu, Ge Bai, Dong-Sheng Wang, Yuexuan Wang and Giulio Chiribella

Deep neural networks are a powerful tool for the characterization of quantum states. Existing networks are typically trained with experimental data gathered from the specific quantum state that needs to be characterized. But is it possible to train a neural network offline and to make predictions about quantum states other than the ones used in the training stage? Here we introduce a model of network that can be trained with classically simulated data from a fiducial set of states and measurements, and can later be used to experimentally characterize quantum states that share structural similarities with the states in the fiducial set. With little guidance of quantum physics, the network builds its own data-driven representation of quantum states, and then uses it to predict the outcome statistics of quantum measurements that have not been performed yet. The state representation produced by the network can also be used for tasks beyond the prediction of outcome statistics, including clustering of states and the identification of different phases of matter. Our network model provides a flexible approach that can be applied to online learning scenarios, where predictions must be generated as soon as experimental data become available, and to blind learning scenarios where the learner has only access to an encrypted description of the quantum hardware.

Predicting Gibbs State Expectation Values with Pure Thermal Shadows

Luuk Coopmans, Yuta Kikuchi and Marcello Benedetti

The preparation and computation of many properties of quantum Gibbs states are essential for algorithms such as quantum semi-definite programming and quantum Boltzmann machines. We propose a quantum algorithm that can predict M linear functions of an arbitrary Gibbs state with only O(log M) experimental measurements. Our main insight is that for sufficiently large systems we do not need to prepare the n-qubit mixed Gibbs state explicitly, but instead we can evolve a random n-qubit pure state in imaginary time. The result then follows by constructing classical shadows of these random pure states. We propose a quantum circuit that implements this algorithm by using quantum signal processing for the imaginary time evolution. We verify the algorithm numerically by simulating the circuit for a 10 spin-1/2 XXZ-Heisenberg model.

Learning quantum channels without input control

Marco Fanizza, Yihui Quek and Matteo Rosati

We introduce a general statistical learning theory for processes that take as input a classical random variable and output a quantum state. This is a practically motivated setting where the learner wants to find an approximate representation of an unknown quantum process, without having control on the classical inputs that govern the process itself, e.g., in learning astronomical processes, Hamiltonian evolution at variable times, and biological processes not controlled by the observer. We provide an algorithm for learning with high probability in this setting with a finite amount of samples. Our results indicate that in learning quantum processes, too, we may break the "curse of dimensionality": an infinite concept class can be learned with a finite number of samples. As a concrete example of processes that can be learned in this manner, we show that our bounds apply to learning the (infinite) class of local random circuits on computational basis states encoding classical strings, and to a generalization of binary functions in the quantum setting. Finally, we obtain as a by-product sufficient conditions for performing shadow tomography of classical-quantum states.

Session: Quantum Learning Methods - Part I

Unifying Quantum Machine Learning Models: Theory and Practical Implications

Sofiene Jerbi, Lukas Fiderer, Hendrik Poulsen Nautrup, Jonas Kübler, Hans Briegel and Vedran Dunjko

We present a unifying framework for all standard quantum machine learning models proposed in the literature. Through this framework, we show that linear quantum models are universal models, in the sense that they can represent (either exactly or approximately) any (quantum) machine learning model. Nonetheless, we also show the existence of learning tasks where data re-uploading models present exponential advantages compared to linear and kernel quantum models in terms of the number of qubits and the number of data points they need to learn.

Pricing multi-asset derivatives by finite difference method on a quantum computer

Koichi Miyamoto, Kenji Kubo, Kosuke Mitarai and Keisuke Fujii

Recently, there are growing interests in applications of quantum computing to finance. In this presentation, we propose the quantum algorithm for derivative pricing based on solving the Black-Scholes partial differential equation by the finite difference method (FDM). Classically, the FDM suffers from the curse of dimensionality, the exponential growth of complexity in the case of multiple underlying assets. Our algorithm provides the exponential speedup compared to the classical FDM, using the quantum algorithm for solving differential equations. Note that the quantum differential equation solver outputs the state $|\mathrm{V}\rangle$ that encodes the solution but we actually want the derivative price V0 for the present underlying asset prices. Naively, extracting V0 from $|\mathrm{V}\rangle$ takes an exponential complexity, which ruins the quantum speedup. To address this point, we calculate V_0 as the inner product of $|\mathrm{V}\rangle$ and the state that encodes the asset price probability distribution. This algorithm is for a fault-tolerant quantum computer, and we also propose another version of the algorithm for a Noisy Intermediate-Scale Quantum computer. In this version, we use the variational quantum simulation to create $|\mathrm{V}\rangle$ and again compute V_0 as the inner product. We conduct a numerical experiment to validate our algorithm.

Learning via Many Body Localized Hidden Born Machine

Sona Najafi, Weishun Zhong, Xun Gao and Susanne Yelin

Generative models lie at the heart of unsupervised learning providing a direct way of estimating joint probability distribution. Inspired by the probabilistic nature of quantum mechanics, Born machine has been risen as promising quantum generative model with capability to learn both classical and quantum data. While tensor networks has been primary focus of the Born Machine, here, based on our recent work [8], we utilize the many body localize hidden Born Machine as a resource of learning. We devise our Born machine based on a sequence of quench governed by MBL Hamiltonian such that by optimizing the disorder values we reach our target. We show that our MBL Bor Machine is capable of learning different category of data such as MNIST pattern. We attribute such learning capability of MBL phase to emergence of memory and it's integrating nature.

Session: Quantum Learning Methods - Part II

Building spatial symmetries into parameterized quantum circuits for faster training

Frederic Sauvage, Martin Larocca, Patrick Coles and Marco Cerezo

Practical success of quantum learning models hinges on having a suitable structure for the parameterized quantum circuit employed. Such structure is defined both by the type of gates chosen and by the correlation of their parameters. While much research has been devoted to devising adequate gatesets, very little is known about how their parameters should be structured. In this work, we show that an ideal parameter structure naturally emerges when carefully considering spatial symmetries (i.e., the symmetries that are permutations of parts of the system under study). Namely, we consider the automorphism group of the problem Hamiltonian, leading us to develop a circuit construction that is equivariant under this symmetry group. The benefits of our methodology are numerically probed in many ground-state problems, revealing a consistent improvement (in terms of circuit depth, number of parameters required, and trainability) compared to current circuit constructions.

Symbiosis of Tensor Networks and Quantum Circuits in Training Parametrized Quantum Algorithms

Manuel S. Rudolph, Atithi Acharya, Jing Chen, Jacob Miller and Alejandro Perdomo-Ortiz

Quantum computing resources are currently sparse, expensive, and challenging to use effectively. To utilize them to their fullest capabilities, it is required that we uncover use-cases where classical computing techniques are not sufficient to solve the task, and then set a quantum algorithm up for success. Most near-term approaches of quantum algorithms employ Parametrized Quantum Circuits (PQCs), which may be able to solve a range of valuable tasks on relatively noisy quantum devices. PQCs are however hindered by unfavorable performance guarantees due to generic quantum circuit design, parameter initialization, and optimization procedures. In contrast, quantum-inspired models based on Tensor Networks (TNs) are emerging as a powerful technology to approximate quantum systems on classical computers, raising the bar for quantum computers to overcome. In our work, we demonstrate a practical joint optimization framework utilizing TNs and PQCs, where the task of interest is first partially solved with more abundant classical resources, and the intermediate solution is then mapped to quantum circuits, where optimization can continue on quantum computers. To fully leverage the classical solution, we present a technique to map Matrix Product States (MPS) of arbitrary bond dimension to PQCs with a low number of layers of parametrizable two-qubit gates. We demonstrate that our framework avoids common pitfalls of randomly initialized parameters, which drastically improves the performance of deep quantum circuits and can convert PQCs that are practically untrainable to ones that reliably converge to high-quality solutions.

Bayesian Learning of Parameterised Quantum Circuits

Samuel Duffield, Marcello Benedetti and Matthias Rosenkranz

Currently available quantum computers suffer from constraints including hardware noise and a limited number of qubits. As such, variational quantum algorithms that utilise a classical optimiser in order to train a parameterised quantum circuit have drawn significant attention for near-term practical applications of quantum technology. In this work, we take a probabilistic point of view and reformulate the classical optimisation as an approximation of a Bayesian posterior. The posterior is induced by combining the cost function to be minimised with a prior distribution over the parameters of the quantum circuit. We describe a dimension reduction strategy based on a maximum a posteriori point estimate with a Laplace prior. Experiments on the Quantinuum H1-2 computer show that the resulting circuits are faster to execute and less noisy than the circuits trained without the dimension reduction strategy. We subsequently describe a posterior sampling strategy based on stochastic gradient Langevin dynamics. Numerical simulations on three different problems show that the strategy is capable of generating samples from the full posterior and avoiding local optima.

Parameterized quantum circuits for rare dynamic learning

Alissa Wilms, Laura Ohff, Andrea Skolik and Jens Eisert

In the study of non-equilibrium or industrial systems, rare events are crucial for understanding the systems' behavior and the effective search for such rare dynamics is frequently the subject of research. Since they are atypical, one requires specific methods for sampling and generating rare event statistics in an automated and statistically meaningful way. Recent publications have shown variational quantum algorithms to be among the most promising candidates for near-term applications on quantum devices. In this article, we propose a quantum reinforcement learning (QRL) approach to study rare dynamics of time-dependent systems and investigate its benefits over classical approaches based on neural networks. We demonstrate that QRL can learn and generate the rare dynamics of a random walker. Furthermore, we are able to show better learning behavior, with fewer parameters in comparison to the classical counterpart. This is the first study of QRL in rare event statistics and suggests that QRL is a viable method to study rare dynamics of a system.

Group-Invariant Quantum Machine Learning

Martin Larocca, Frederic Sauvage, Faris Sbahi, Guillaume Verdon, Patrick Coles and Marco Vinicio Sebastian de la Roca

Quantum Machine Learning (QML) models are aimed at learning from data encoded in quantum states. Recently, it has been shown that models with little to no inductive biases (i.e., with no assumptions about the problem embedded in the model) are likely to have trainability and generalization issues, especially for large problem sizes. As such, it is fundamental to develop schemes that encode as much information as available about the problem at hand. In this work we present a simple, yet powerful, framework where the underlying invariances in the data and task are used to build QML models that, by construction, respect those symmetries. Specifically, these group-invariant models produce outputs that remain fixed under the action of some symmetry group G associated with the learning task. We first present theoretical results underpinning the design of G-invariant models, and then exemplify their application through several paradigmatic QML classification tasks. Notably, our framework allows us to recover, in an elegant way, several well known algorithms for the literature, as well as to discover new ones. Taken together, our results pave the way towards a more effective QML model design.

Invited Talk: Shahnawaz Ahmed

Chalmers University of Technology, Sweden

Machine learning quantum states and operations: from neural networks to optimization on manifolds

Machine learning techniques have found recent applications in quantum tomography. The underlying idea is to use an efficient ansatz to represent a quantum state or process and learn it from data. Neural network architectures from Restricted Boltzmann Machines to Recurrent Neural Networks have been proposed as ansatzes for quantum states. Such ansatzes can be trained using standard gradient-based optimization to directly estimate a quantum state's density matrix or allow efficient sampling of measurement outcomes. Similar ideas can be applied to learn a quantum process. In this talk, we will discuss several such machine learning methods for quantum state and process tomography. We will elucidate the necessary ingredients to apply machine learning to the tomography problem - from using physics-based constraints on the ansatzes to constraints on the training itself such as gradient-descent on a manifold. We will also compare machine learning to existing standard techniques such as maximum likelihood estimation, compressed sensing or projection-based algorithms to show how ideas from machine learning can enhance the set of tools for quantum characterization.

Invited Talk: Mario Krenn

Max Planck Institute, Germany

Towards an Artificial Muse for new Ideas in Quantum Physics

Articial intelligence (AI) is a potentially disruptive tool for physics and science in general. One crucial question is how this technology can contribute at a conceptual level to help acquire new scientic understanding or inspire new surprising ideas. I will talk about how AI can be used as an articial muse in quantum physics, which suggests surprising and unconventional ideas and techniques that the human scientist can interpret, understand and generalize. [1] Krenn, Kottmann, Tischler, Aspuru-Guzik, Conceptual understanding through efficient automated design of quantum optical experiments. Physical Review X 11(3), 031044 (2021). [2] Krenn, Pollice, Guo, Aldeghi, Cervera-Lierta, Friederich, Gomes, Hase, Jinich, Nigam, Yao, Aspuru-Guzik, On scientic understanding with articial intelligence. arXiv:2204.01467 (2022). [3] Krenn, Zeilinger, Predicting research trends with semantic and neural networks with an application in quantum physics. PNAS 117(4), 1910-1916 (2020).

Session: Quantum Learning Methods - Part III

Quantum Reinforcement Learning via Policy Iteration

El Amine Cherrat, Iordanis Kerenidis and Anupam Prakash

We provide a general framework for performing quantum reinforcement learning via policy iteration in [3]. We validate our framework by designing and analyzing: quantum policy evaluation methods for infinite horizon discounted problems by building quantum states that approximately encode the value function of a policy π; and quantum policy improvement methods by post-processing measurement outcomes on these quantum states. Last, we study the theoretical and experimental performance of our quantum algorithms on two environments from gym.

Quantum Extremal Learning

Savvas Varsamopoulos, Evan Philip, Herman W. T. van Vlijmen, Sairam Menon, Ann Vos, Natalia Dyubankova, Bert Torfs, Anthony Rowe and Vincent E. Elfving

We introduce the concept of quantum extremal learning(QEL), which is the process of finding the input to a hidden function that extremizes the function output, without having direct access to the hidden function, given only partial input-output training data. We describe a QEL algorithm that consists of a parametric quantum circuit that is variationally trained to model data input-output relationships and a trainable quantum feature map. The feature map encodes the input data to a qubit register and is differentiated to find the input that extremizes the model output. Both modeling and optimization can be run on the same circuit/quantum chip based on established (quantum) machine learning processes. In this manuscript, we will showcase results on Max-Cut problems where discrete classical data are assumed. We envision that due to its general framework and simple construction, the QEL algorithm will be able to solve a wide variety of applications in different fields.

Quantum versus Classical Autoencoders

Coral Featherstone and Francesco Petruccione

An autoencoder is a neural network that stores a compressed representation of the input passed to it in a manner that captures the essential features of the data, and discards the detail, with the restriction that it doesn't reproduce exact copies of the input. Autoencoders are designed to produce approximations of the data model they are learning. Undercomplete autoencoders facilitate compression by starting and ending a neural network with a particular layer size and passing through a layer containing less neurons than the input and output layers, forcing the network to discard non-essential features. The undercomplete autoencoder appears in the quantum literature in three variants. Other classical autoencoders retain the charactersitics required of an autoencoder; but, don't do so in the typical bottleneck configuration. These variants do not appear in the quantum autoencoder literature. The aim of this paper is to describe several of the classical autoencoder variants, and to compare them to quantum autoencoders.

A Quantum Enhanced Learning Algorithm for Maze Problems

Oliver Sefrin and Sabine Wölk

In reinforcement learning, a so-called agent should learn to optimally solve a given task by performing actions within an environment. As an example, we consider the grid-world, a two-dimensional maze for which the shortest way from an initial position to a given goal has to be found. The agent receives rewards for helpful actions which enables him to learn optimal solutions. For large action spaces, a mapping of actions to a quantum setting can be beneficial in finding rewarded actions faster and thus in speeding up the learning process. This speed-up can be achieved by oracularizing the environment and performing amplitude amplification. Based on this technique, a hybrid agent which alternates between quantum and classical behavior has been developed previously for deterministic and strictly epochal environments. Here, strictly epochal means that an epoch consists of a fixed number of actions, after which the environment is reset to its initial state. We present and analyze strategies which aim at resolving the hybrid agent's current restriction of searching for action sequences with a fixed length. This is a first step towards applying the hybrid agent on environments with a generally unknown optimal action sequence length such as in the grid-world problem.

Hyperparameter Importance of Quantum Neural Networks Across Small Datasets

Charles Moussa, Jan N. van Rijn, Thomas Baeck and Vedran Dunjko

As restricted quantum computers are slowly becoming a reality, the search for meaningful first applications intensifies. In this domain, one of the more investigated approaches is the use of a special type of quantum circuit - a so-called quantum neural network – to serve as a basis for a machine learning model. Roughly speaking, as the name suggests, a quantum neural network can play a similar role to a neural network. However, specifically for applications in machine learning contexts, very little is known about suitable circuit architectures, or model hyperparameters one should use to achieve good learning performance. In this work, we apply the functional ANOVA framework to quantum neural networks to analyze which of the hyperparameters were most influential for their predictive performance. We analyze one of the most typically used quantum neural network architectures. We then apply this to 7 open-source datasets from the OpenML-CC18 classification benchmark whose number of features is small enough to fit on quantum hardware with less than 20 qubits. Three main levels of importance were detected from the ranking of hyperparameters obtained with functional ANOVA. Our experiment both confirmed expected patterns and revealed new insights. For instance, setting well the learning rate is deemed the most critical hyperparameter in terms of marginal contribution on all datasets, whereas the particular choice of entangling gates used is considered the least important except on one dataset. This work introduces new methodologies to study quantum machine learning models and provides new insights toward quantum model selection.

Multiclass classification based on quantum state discrimination

Giuseppe Sergioli, Andres Camilo Granda Arando, Roberto Giuntini, Hector Carlos Freytes and Federico Hernan Holik

We address the problem of multiclass classification using quantum-inspired algorithms which are based on quantum state discrimination techniques. The most important feature of these algorithms is that they can be implemented in a classical hardware, still offering advantages. We provide a comparison between the multiclass quantum-inspired classifier with other standard algorithms based in the one vs one and one vs rest strategies. We also show a quantum implementation of the test part, showing the differences with its classical counterpart and posing some questions for future research.

Session: Computational Learning Theory

Clifford Circuits can be Properly PAC Learned if and only if RP=NP

Daniel Liang

Given a dataset of input states, measurements, and probabilities, is it possible to efficiently predict the measurement probabilities associated with a quantum circuit? Recent work of Caro and Datta [19] studied the problem of PAC learning quantum circuits in an information theoretic sense, leaving open questions of computational efficiency. In particular, one candidate class of circuits for which an efficient learner might have been possible was that of Clifford circuits, since the corresponding set of states generated by such circuits, called stabilizer states, are known to be efficiently PAC learnable [37]. Here we provide a negative result, showing that proper learning of CNOT circuits is hard for classical learners unless $RP = NP$. As the classical analogue and subset of Clifford circuits, this naturally leads to a hardness result for Clifford circuits as well. Additionally, we show that if $RP = NP$ then there would exist efficient proper learning algorithms for CNOT and Clifford circuits. By similar arguments, we also find that an efficient proper quantum learner for such circuits exists if and only if $NP \subseteq RQP$.

A single T-gate makes distribution learning hard

Marcel Hinsche, Marios Ioannou, Alexander Nietner, Jonas Haferkamp, Yihui Quek, Dominik Hangleiter, Jean-Pierre Seifert, Jens Eisert and Ryan Sweke

In this work, we provide an extensive characterization of the learnability of the output distributions of local quantum circuits. Our first result yields insight into the relationship between the efficient learnability and the efficient simulatability of these distributions. Specifically, we prove that the density modelling problem associated with Clifford circuits can be efficiently solved, while for depth $d = n^{\Omega(1)}$ circuits the injection of a single T-gate into the circuit renders this problem hard. Our second set of results provides insight into the potential and limitations of quantum generative modelling algorithms. We first show that the generative modelling problem associated with depth $d = n^{\Omega(1)}$ local quantum circuits is hard for any learning algorithm, classical or quantum. As a consequence, one cannot use a quantum algorithm to gain a practical advantage for this task. We then show that, for a wide variety of the most practically relevant learning algorithms – including hybrid-quantum classical algorithms – even the generative modelling problem associated with depth $d = \omega(\log(n))$ Clifford circuits is hard. This result places limitations on the applicability of near-term hybrid quantum-classical generative modelling algorithms.

Classical surrogates for quantum learning models

Franz J. Schreiber, Jens Eisert and Johannes Jakob Meyer

The advent of noisy intermediate-scale quantum computers has put the search for possible applications to the forefront of quantum information science. One area where hopes for an advantage through near-term quantum computers are high is quantum machine learning, where variational quantum learning models based on parametrized quantum circuits are discussed. In this work, we introduce the concept of a classical surrogate, a classical model which can be efficiently obtained from a trained quantum learning model and reproduces its input-output relations. As inference can be performed classically, the existence of a classical surrogate greatly enhances the applicability of a quantum learning strategy. However, the classical surrogate also challenges possible advantages of quantum schemes. As it is possible to directly optimize the ansatz of the classical surrogate, they create a natural benchmark the quantum model has to outperform. We show that large classes of well-analyzed re-uploading models have a classical surrogate. We conducted numerical experiments and found that these quantum models show no advantage in performance or trainability in the problems we analyze. This leaves only generalization capability as possible point of quantum advantage and emphasizes the dire need for a better understanding of inductive biases of quantum learning models.

Exponential concentration and untrainability in quantum kernel methods

Supanut Thanasilp, Samson Wang, Marco Cerezo and Zoe Holmes

Kernel methods in Quantum Machine Learning (QML) have recently gained significant attention as a potential candidate for achieving a quantum advantage in data analysis [1–7]. While the ultimate goal is to obtain a trained model that generalizes well, this crucially depends on how well the model can be trained in the first place. Among other attractive properties, kernel methods are often said to enjoy trainability guarantees due to the convexity of their loss landscapes [8–11]. This is in contrast to Quantum Neural Networks (QNNs) where the loss landscape is generally non-convex [12, 13] and can exhibit barren plateaus (BPs), where cost gradients vanish exponentially in the number of qubits [14–26]. However, the favorability of the kernel-based trainability over QNNs relies heavily on the key assumption that values of the quantum kernel can be efficiently evaluated on a quantum computer. In this work, we study the trainability of quantum kernels in the context of the resource scaling required to accurately estimate the kernel entries values. We show that under certain conditions, the value of quantum kernels can exponentially concentrate (in the number of qubits) around a fixed value. In such cases, the number of shots required to evaluate the kernels to a sufficiently high precision scales exponentially. This indicates that the efficient evaluation of quantum kernels cannot always be taken for granted. The exponential concentration of quantum kernels may be viewed as stemming from the difficulty of extracting information from quantum states. In particular, we identify four key features of a quantum kernel method that can severely hinder the information extraction process. These include: i. the expressibility of data embedding, ii. entanglement, iii. global measurements and iv. noise (see Fig. 1). For each source, we derive an associated concentration bound and support our analytic results with numerical demonstrations for standard datasets and a common embedding scheme. Thus we show that, in contrast to common belief, quantum kernels are not always guaranteed to train successfully and additional care is needed. Altogether, our results serve as a guideline on what features to be avoided in order to ensure the efficient evaluation and training of quantum kernels.

Subtleties in the trainability of quantum machine learning models

Supanut Thanasilp, Samson Wang, Nhat Anh Nghiem Vu, Patrick Coles and Marco Vinicio Sebastian de la Roca

Quantum Machine Learning (QML), has received significant attention due to its promise of making practical use of quantum computers and achieving a speedup over traditional machine learning for data analysis [2– 8]. Its success usually hinges on efficiently training the parameters in Quantum Neural Networks (QNN) to achieve not only small loss function values on the training data but also small generalization error on the previously unseen data [9–12]. However, despite tremendous recent progress in understanding the capabilities of QML models, the field is still lacking rigorous theoretical results for the trainability and gradient scaling of loss functions. Some trainability results have been proven for the closely related field of Variational Quantum Algorithms (VQAs). Here, one also optimizes parametrized quantum circuits to solve some practical useful tasks [13]. Here, a great deal of effort has been put forward towards analyzing and avoiding the barren plateau phenomenon [14–23]. When a VQA exhibits a barren plateau, the cost function gradients vanish exponentially with the problem size (i.e. number of qubits), leading to an exponentially flat optimization landscape. Barren plateaus greatly impact the trainability of the parameters as it is impossible to navigate the flat landscape without expending an exponential amount of resources [24–26]. While both fields involve training a parametrized quantum circuit, there are crucial differences that make the trainability results for VQA setting not readily applicable to QML. This is due to the fact that QML models are more complex in general. For instance, in QML one needs to deal with datasets [27], which further require the use of an embedding scheme when dealing with classical data (i.e., the map between classical data $\mathbf{x}_i$ and quantum states ρ_i) [28–31]. In addition, QML loss functions can be more complicated than VQA cost functions, as the latter are usually linear functions of expectation values of some set of operators. Thus, given the large body of literature studying barren plateaus in VQAs, the natural question that arises is: Are the gradient scaling and barren plateau results also applicable to QML settings? In this work we bridge the two frameworks and show that gradient scaling results for VQAs can also be applied to study the gradient scaling of QML models. Our results indicate that features deemed detrimental for VQA trainability can also lead to issues such as barren plateaus in QML. Consequently, our work shows that many QML proposals in the literature need to be revised. In addition, we provide theoretical and numerical evidence that QML models exhibit further trainability issues not present in VQAs, arising from the use of a training dataset. We refer to these as dataset-induced barren plateaus. These results are most relevant when dealing with classical data, as here the choice of embedding scheme (i.e., the map between classical and quantum data) can greatly affect the gradient scaling.

Quantum Boosting using Domain-Partitioning Hypotheses

Sagnik Chatterjee, Rohan Bhatia, Parmeet Singh Chani and Debajyoti Bera

Boosting is an ensemble learning method that converts a weak learner into a strong learner in the PAC learning framework. Freund and Schapire designed the Godel prize-winning algorithm named AdaBoost that can boost learners which output binary hypotheses. Recently, Arunachalam and Maity presented the first quantum boosting algorithm with similar theoretical guarantees. Their algorithm, which we refer to as QAdaBoost henceforth, is a quantum adaptation of AdaBoost, and only works for the binary hypothesis case. QAdaBoost is quadratically faster than AdaBoost in terms of the VC-dimension of the hypothesis class of the weak learner but polynomially worse in the bias of the weak learner. Izdebski et al. posed an open question on whether we can boost quantum weak learners that output non-binary hypothesis. In this work, we address this open question by developing the QRealBoost algorithm which was motivated by the classical RealBoost algorithm. The main technical challenge was to provide provable guarantees for convergence, generalization bounds, and quantum speedup, given that quantum subroutines are noisy and probabilistic. We prove that QRealBoost retains the quadratic speedup of QAdaBoost over AdaBoost and further achieves a polynomial speedup over QAdaBoost in terms of both the bias of the learner and the time taken by the learner to learn the target concept class. Finally, we perform empirical evaluations on QRealBoost and report encouraging observations on quantum simulators by benchmarking the convergence performance of QRealBoost against QAdaBoost, AdaBoost, and RealBoost on a subset of the MNIST dataset and Breast Cancer Wisconsin dataset.

Invited Talk: Sabre Kais

Purdue University, USA

Information scrambling for navigating the learning landscape of a quantum machine learning model

In this talk, I will focus on quantum machine learning, particularly the Restricted Boltzmann Machine (RBM), as it emerged to be a promising alternative approach leveraging the power of quantum computers. The workhorse of our technique is a shallow neural network encoding the desired state of the system with the amplitude computed by sampling the Gibbs oltzmann distribution using a quantum circuit and the phase information obtained classically from the nonlinear activation of a separate set of neurons. In Addition to present the successful applications for electronic structure of two-dimensional materials I will discuss and illustrate that the imaginary components of out-of-time correlators can be related to conventional measures of correlation like mutual information. Such an analysis offers important insights into the training dynamics by unraveling how quantum information is scrambled through such a network introducing correlation among its constituent sub-systems. This approach not only demystifies the training of quantum machine learning models but can also explicate the capacitive quality of the model.

Keynote Talk: Hartmut Neven

Google Quantum AI, USA

The Google Quantum AI team's current Perspective on Quantum Machine Learning

I will start the talk with an update on the status of the Quantum AI project. Google's Quantum AI team is making steady progress on its roadmap to build a quantum computer with 1000 long lived logical qubits using superconducting electronics. In summer 2022, we demonstrated for the first time that by scaling up a surface code qubit we can reduce its logical error rate. We are now turning our attention to fabricating a long-lived logical qubit with 10^{-6} error rate. Then I will discuss two findings relevant to the quantum machine learning community. A negative result is that quadratic speedups, for example for quantum enhanced optimization, are not sufficient to achieve practically relevant quantum advantage. A positive result is the discovery of quantum machine learning algorithms that can learn properties of quantum systems with exponentially reduced sample complexity.

Abstracts of talks (workshop)

Mingsheng Ying

Institute of Software, Chinese Academy of Science and Tsinghua University, China

Quantum Hoare logic and its applications

Hoare logic is a foundation of axiomatic semantics of classical programs and it provides effective proof techniques for reasoning about correctness of classical programs. To offer similar techniques for quantum program verification and to build a logical foundation of programming methodology for quantum computers, we develop a full-fledged Hoare logic for both partial and total correctness of quantum programs. It is proved that this logic is (relatively) complete by exploiting the power of weakest preconditions and weakest liberal preconditions for quantum programs.

Urs Schreiber

Center for Quantum and Topological Systems - NYU Abu Dhabi, UK

Quantum Data Types via Linear Homotopy Type Theory

The proper concept of data types in quantum programming languages, hence of their formal verification and categorical semantics, has remained somewhat elusive, as witnessed by the issue of "dynamic lifting" encountered in the Quipper language family. In this talk I explain our claim that a powerful quantum data type system elegantly solving these problems is naturally provided by the "linear homotopy type theory" recently realized by M. Riley. Together with our previous claim that homotopy type theory natively knows about the fine detail of su(2)-anyon braid gates, this shows that linear homotopy type theory is a natural substrate for typed quantum programming languages aware of topological quantum hardware.
This is joint work, at CQTS, with D. J. Myers, M. Riley and H. Sati.
Notes are available at ncatlab.org/schreiber/show/TQPinLHoTT

Kevin Mato

Technical University of Munich, Dept. of Electrical and Computer Engineering, Germany

MQT: The Munich Quantum Toolkit

Quantum computers are approaching the turning point of practical applicability, and there is a thriving community that creates toolkits to provide design automation methods. Having efficient software solutions for design automation will be an enabling factor for quantum computing, just as design automation was an enabler for classical computers in the past decades. In order to tackle the challenges of leveraging quantum technologies, we require new and dedicated tooling; however, the experience from classical design can help to overcome the complexity of the new paradigms. We present the Munich Quantum Toolkit (MQT), a set of tools for quantum computing developed at the Technical University of Munich that utilizes this design automation expertise. MQT offers complementary approaches for many design problems in quantum computing such as simulation, compilation, and verification. In the following, we provide an introduction to design automation and to the toolkit, for potential users and developers.

Ilmo Salmenperä and Jukka Nurminen

University of Helsinki, Finland

Software Solutions to Hardware Limitations for Quantum Sampled Restricted Boltzmann Machines

Restricted Boltzmann Machines are common machine learn- ing models that can utilise quantum annealing devices in their training processes as quantum samplers. While this approach has shown promise as an alternative for classical sampling methods, limitations of quantum hardware, such as the amount of qubits and the lack of connectivity be- tween the qubits, still poses as a barrier for wide scale adoption. The key contributions of this study is evaluation of software techniques such as unit dropout method, passive labelling and parallelisation techniques for addressing these hardware limitations. The study found that quantum sampling showed comparable results to its classical counterparts in certain contexts, while in others the performance suffered due to the difficulty of evaluating the device parameters of the quantum annealer.

Johanna Barzen

University of Stuttgart, Germany

QHAna: A tool for comparing classical and quantum machine learning

The development of quantum applications usually involves the development of quantum programs, classical programs, and workflows to orchestrate them. Therefore, several lifecycles of these software artifacts must be integrated into an overall quantum software development lifecycle. A key issue in such a lifecycle is deciding when to choose classical and when to choose quantum-based programs. In machine learning in particular, several quantum-based algorithms already exist, but deciding which of these algorithms are useful and provide meaningful results compared to their classical counterparts is time-consuming and requires a deep understanding of the underlying algorithms. This is where our tool QHAna chimes in: it provides support for comparing the results of classical and quantum-based machine learning, based on a plugin architecture that can be easily customized for different use cases to enable initial application knowledge.

Sebastian Feld

QuTech, Delft, Netherlands

Machine Learning for Quantum System Software

As with classical computers, also quantum computers need system software to be operated. Since we are currently in the so-called NISQ-era and thus quantum devices are small and prone to error, developers need to take care of a plethora of problems in order to make the quantum devices usable. This talk highlights major possibilities to utilize machine learning to create efficient, scalable, and future-proof system software for quantum computing devices.

Medina Bandic

TU Delft, Netherlands

Characterizing quantum benchmarks for improved mapping of quantum circuits

Quantum circuit mapping techniques are crucial for successfully executing quantum algorithms on current resource-constrained and error-prone quantum processors. They perform some modifications on the quantum circuit in a way that it complies with the hardware restrictions, while trying to minimize the resulting gate and depth overhead to in turn increase the circuit success rate. Most quantum circuit mapping techniques focus on the hardware properties, though some works have already pointed out the importance of also considering algorithm characteristics. In this talk the focus is on how we performed a thorough characterization of quantum circuits by not only extracting typical parameters like the number of qubits and gates, but also graph theory-based metrics from their corresponding qubit interaction and gate-dependency graphs, as well as on how can we use this to further on improve the quantum circuit mapping process.

Nicola Assolini

Università di Verona, Italy

A Static Analysis of Uncomuptation

There are several important aspects of quantum programming that differ from conventional programming. As an example, while analogously to the classical setting, quantum computations often produce temporary values, removing such values is much more challenging in quantum computation as it induces an implicit measurement collapsing the state and producing unintended side-effects due to entanglement. We define an analysis of uncomputation that allows us to safely discharging temporary variables in a quantum program. To this purpose we exploit data-flow analysis techniques, which are typically employed in classical compilers, for statically detecting variables that need to be uncomputed.

Koen Bertels

University of Porto, Portugal

We need more PISQ-driven Quantum Logic Research

Classical hardware will continue to be present in almost all aspects of industry, science, and daily life. All computer chips are based on transistors, and we are very close to the end of making them smaller than 2nm. Quantum computing is considered by many people and companies world-wide as one of the alternatives for building completely new and even more powerful computers. This technology choice will have serious impacts on basically all scientific fields and industrial products that can be made. We should be aware that quantum logic is substantially different than any classical way of logical thinking and that it takes around 10 to 15 years before we have formulated all the necessary concepts for any scientific field. And that a new generation of quantum-competent people has been trained and can start working on concrete problems. The quantum physical research for producing good quality qubits still need serious improvements and that is world-wide called Noisy Intermediate Scale Quantum based research, abbreviated to NISQ. But it is becoming very urgent that scientists in any scientific field and from all over the world need to start doing research in their domain but using quantum logic concepts. This kind of research can be called PISQ, which stands for Perfect (Qubits) Intermediate Scale Quantum. Within QBee and driven by PhD-level scientists, we have started looking at quantum computational chemistry and mage classification of satellite images taken from our planet. This presentation will briefly discuss those application fields as well as the international collaboration needed to achieve a productive and scientific agreement in all fields.

Registration: Hall, Centro Congressi Federido II, via Partenope 36, 80121 Naples, Italy

Talks: Aula Magna, Centro Congressi Federido II, via Partenope 36, 80121 Naples, Italy

Posters: Sala A1, Centro Congressi Federido II, via Partenope 36, 80121 Naples, Italy

Welcome Reception: Pizzeria Salvo, Riviera di Chiaia 271, 80121 Naples, Italy

Social Dinner: Pulcinella Dining Room, Hotel Royal Continental, via Partenope 38/44, Lungomare Caracciolo, 80121 Naples, Italy