

Quantum Physics and Machine Learning



Universitätszentrum Obergurgl

March 24th - April 1st. 2024

Learning unitaries with quantum statistical queries

Armando Angrisani

EPFL

We propose several algorithms for learning unitary operators from quantum statistical queries (QSQs) with respect to their Choi-Jamiolkowski state. Quantum statistical queries capture the capabilities of a learner with limited quantum resources, which receives as input only noisy estimates of expected values of measurements.

Our methods hinge on a novel technique for estimating the Fourier mass of a unitary on a subset of Pauli strings with a single quantum statistical query, generalizing a previous result for uniform quantum examples. Exploiting this insight, we show that the quantum Goldreich-Levin algorithm can be implemented with quantum statistical queries, whereas the prior version of the algorithm involves oracle access to the unitary and its inverse. Moreover, we prove that $O(\log(n))$ -juntas and quantum Boolean functions with constant total influence are efficiently learnable in our model, and constant-depth circuits are learnable sample-efficiently with quantum statistical queries. On the other hand, all previous algorithms for these tasks require direct access to the Choi-Jamiolkowski state or oracle access to the unitary.

We also demonstrate that, despite these positive results, quantum statistical queries lead to an exponentially larger sample complexity for certain tasks, compared to separable measurements to the Choi-Jamiolkowski state. In particular, we show an exponential lower bound for learning a class of phase-oracle unitaries and a double exponential lower bound for testing the unitarity of channels, adapting to our setting previous arguments for quantum states.

Machine learning for quantum device control

Natalia Ares

University of Oxford

Machine learning is proving to be essential in the tuning and characterization of quantum devices. The search for operation conditions, which often requires navigating large and complex parameter spaces, can now be fully automated, with performances superior to those achieved by human experts. Now these machine learning approaches are not only enabling scalability by automating qubit control, but also by providing us with unprecedented insight into quantum device variability.

We can use these machine learning algorithms for automatic tuning across different semiconductor platforms. This demonstrates not only the robustness of these algorithms against the differences in the characteristics of the material system and device architecture, but that they can provide a tool for their comparison and analysis. I will show that by using a physics-aware machine learning algorithm we are able to infer the disorder potential affecting the operation of quantum dot devices, revealing a hidden characteristic of such devices, and thus narrowing the gap between simulation and reality.

Expressivity of parameterized quantum circuits for generative modeling of continuous multivariate distributions

Alice Barthe

CERN / Leiden University

Parameterized quantum circuits have been extensively used as the basis for machine learning models in regression, classification, and generative tasks. For supervised learning their expressivity has been thoroughly investigated and several universality properties have been proven. However, in the case of quantum generative modeling, the situation is less clear, especially when the task is to model distributions over continuous variables.

In this work, we focus on expectation value sampling-based models; models where random variables are sampled classically, encoded with a parametrized quantum circuit, and the expectation value of fixed observables is measured and returned as a sample. We prove the universality of such variational quantum algorithms for the generation of multivariate distributions. Additionally, we provide a detailed analysis of these models, including fundamental upper bounds on the dimensionality of the distributions these models can represent. We further present a tight trade-off result connecting the needed number of measurements and qubit numbers in order to have universality for a desired dimension of output distribution within an error tolerance.

Finally we also show that the data encoding strategy relates to the so-called polynomial chaos expansion, which is an analog of the Fourier expansion. Our results may help guide the design of future quantum circuits in generative modeling tasks.

Thermodynamics Overfitting and Generalization: Energetic Limits on Predictive Complexity

Alexander Boyd

Trinity College Dublin

Efficiently harvesting thermodynamic resources requires a precise understanding of their structure. This becomes explicit through the lens of information engines, which use information as fuel. Maximizing the work harvested from information corresponds to machine learning, driving information engines to develop complex predictive memory that senses temporal correlations in their environment. This process is known as Thermodynamic Machine Learning, and it is formally equivalent to Maximum-Likelihood Estimation.

We show that an information engine's complex predictive memory poses both energetic benefits and risks. While more memory allows for detection of hidden patterns in information, it also opens the possibility of Thermodynamic Overfitting, where the engine dissipates energy in validation. To address overfitting, we introduce Thermodynamic Regularizers, which incur a cost to engine complexity in training due to the physical constraints on the information engine. We see that regularized Thermodynamic Machine Learning effectively generalizes. The physical constraints from which regularizers are derived improve the effectiveness of learning, adding credence to the view that the laws of physics conspire to create the conditions for emergent intelligence.

Understanding quantum machine learning also requires rethinking generalization

Carlos Bravo-Prieto

FU Berlin

Quantum machine learning models have shown successful generalization performance even when trained with few data. In this talk, we will show that traditional approaches to understanding generalization fail to explain the behavior of such quantum models. We experimentally reveal that state-of-the-art quantum neural networks accurately fit random states and random labeling of training data. This ability to memorize random data defies current notions of small generalization error, problematizing approaches that build on complexity measures such as the VC dimension, the Rademacher complexity, and all their uniform relatives. We complement the empirical results with a theoretical construction showing that quantum neural networks can fit arbitrary labels to quantum states, hinting at their memorization ability. Our results do not preclude the possibility of good generalization with few training data but rather rule out any possible guarantees based only on the properties of the model family. These findings expose a fundamental challenge in the conventional understanding of generalization in quantum machine learning and highlight the need for a paradigm shift in the design of quantum models for machine learning task

Transdisciplinary research for AI in science

Andrew Briggs

University of Oxford

The application of machine learning to quantum physics has enormous potential, both for running experiments efficiently and for discovering new correlations. To fully benefit from AI in the laboratory we need to consider: (i) what are the distinctive attributes of human thought processes, and which of those can be better done by machines; (ii) what distinguishes beneficial uses of AI from harmful uses? [1] Answering these questions requires the kind of transdisciplinary research that is used for addressing wicked societal challenges where the stakes are high, the facts are uncertain, and the values are in dispute, but the decisions are nevertheless urgent.

[1] Andrew Briggs and Michael J. Reiss, *Human Flourishing*, Oxford University Press (2021)

Learning-based quantum error mitigation with near-Clifford circuits

Piotr Czarnik

Jagiellonian University

I will present an error mitigation scheme learning an ansatz to correct observables and circuits similar to classically simulable observables for near-Clifford circuits used as a training set [1]. I will discuss its generalizations using multiple noise levels and Virtual Distillation data to learn an extrapolation to the zero noise limit [2,3]. Furthermore, I will outline a shot-efficient approach to generate the training data [4] and a method to optimize the robustness of the mitigated results [5].

[1] P. Czarnik, A. Arrasmith, P. J. Coles, L. Cincio, Error mitigation with Clifford quantum-circuit data, *Quantum* 5, 592 (2021)

[2] A. Lowe, M. H. Gordon, P. Czarnik, A. Arrasmith, P. J. Coles, L. Cincio, Unified approach to data-driven quantum error mitigation, *Physical Review Research* 3 (3), 033098 (2021)

[3] D. Bultrini, M. H. Gordon, P. Czarnik, A. Arrasmith, M. Cerezo, P. J. Coles, L. Cincio, Unifying and benchmarking state-of-the-art quantum error mitigation techniques, *Quantum* 7, 1034 (2023)

[4] P. Czarnik, M. McKerns, A. T. Sornborger, L. Cincio, Improving the efficiency of learning-based error mitigation, [arXiv:2204.07109](https://arxiv.org/abs/2204.07109)

[5] P. Czarnik, M. McKerns, A. T. Sornborger, L. Cincio, Robust design under uncertainty in quantum error mitigation, [arXiv:2307.05302](https://arxiv.org/abs/2307.05302)

Intepretable machine learning of phases of matter

Anna Dawid

Flatiron Institute

Recently, machine learning has become a powerful tool for detecting quantum phases. While the sole information about the presence of transition is valuable, the lack of interpretability and knowledge on the detected order parameter prevents this tool from becoming a customary element of a physicist's toolbox. Here, we report designing a special convolutional neural network with adaptive kernels, which allows for fully interpretable and unsupervised detection of local order parameters out of spin configurations measured in arbitrary bases. With the proposed architecture, which we call TetrisCNN, we detect relevant and simplest order parameters for the one-dimensional transverse-field Ising model from any combination of projective measurements in the x , y or z basis. More importantly, we tackle experimental measurements of the square Ising and XY models and extract relevant order parameters. We are also considering extending the proposed approach to different lattice geometries and detecting topological or nematic order parameters. This work can lead to integrating ML methods with quantum simulators studying new exotic phases of matter.

Hyperparameter Exploration in Quantum Neural Networks

Vincenzo De Maio

TU Wien

The advent of the Post-Moore era is driving the High Performance Computing (HPC) research community towards alternative architectures, also known as Post-Moore computing. Quantum computing is particularly appealing for HPC, due to proven theoretical speedups for certain problems. A promising candidate is Quantum Machine Learning (QML), where quantum phenomena are exploited for speedups or improved predictive performance. However, different challenges hinder the adoption of QML: firstly, classical data has to be encoded into a quantum state. Considering the current limitations of quantum machines, inefficient encoding can affect potential speedups. Also, near-term QML is usually based on Variational Quantum Algorithms, which are subject to trainability issues, which can, for example, be mitigated by carefully choosing the architecture and initial parameters, i.e., the optimizer, the initial point, and the ansatz. In this talk, we summarize the work of our group in optimizing hyperparameters for QML models on four real-world datasets. We provide insights for the research community on tuning QML models to improve accuracy and reduce training and inference time of the models.

Analogue Quantum Simulators - table-top experiments for quantum field theory

Sebastian Erne

Atominstitut, TU Wien

Analogue (quantum) simulators have gained increasing traction over the last decade as an interdisciplinary branch between experimental and theoretical physics. They present an ideal platform to study fundamental questions of quantum field theory, which are otherwise not or only indirectly accessible in experiments. After giving an introduction to analogue gravity and cosmology based on classical and quantum fluids, I will present recent results of our research on analogue QFT simulators in and beyond the linear regime of the sine-Gordon model. The aim of this talk is to give an overview of the general ideas behind continuous (quantum) field theory simulators based on ultracold atom systems, and the opportunities these offer for the experimental exploration of QFT in flat and curved spacetime and the future directions of physics education and research in general.

Quantum state and process tomography with machine learning and gradient descent

Anton Frisk Kockum

Chalmers University of Technology

The ability to quickly and accurately characterise quantum states and dynamics is crucial for the development of quantum technologies. However, the problem of learning a general quantum state or process has exponential complexity in the size of the quantum system. In this talk, I will present some recent progress we have made for both quantum state and process tomography. For state tomography, I will show how generative adversarial neural networks can outperform standard methods in terms of both the amount of time and data needed [1,2]. For process tomography, I will show how optimization using constrained gradient descent can work both for instances with little data and for larger systems, regimes which previously required two different methods [3]. Our codes for both state and process tomography are freely available on GitHub [4].

[1] S. Ahmed et al., Phys. Rev. Lett. 127, 140502 (2021)

[2] S. Ahmed et al., Phys. Rev. Res. 3, 033278 (2021)

[3] S. Ahmed et al., Phys. Rev. Lett. 130, 150402 (2023)

[4] <https://github.com/quantshah/qst-cgan>; <https://github.com/quantshah/gd-qpt>

Fast and accurate decoding of stabilizer codes for quantum error correction using graph neural networks

Mats Granath

University of Gothenburg

To leverage the full potential of quantum error-correcting stabilizer codes it is crucial to have an efficient and accurate decoder. Accurate, maximum likelihood, decoders are computationally very expensive whereas more efficient decoders give sub-optimal performance. In addition, the accuracy will generally depend on the quality of models and estimates of relevant error rates. I will give an introduction to quantum error correction using stabilizer codes, and discuss our work [Lange et al. arXiv:2307.01241] on a model-free, data-driven, approach to decoding, using a graph neural network (GNN). The decoding problem is formulated as a graph classification task to predict the most likely logical error class. We show that the GNN-based decoder can outperform a matching decoder for circuit level noise on the surface code given only simulated experimental data. We also find that we can use real experimental data [Google Quantum AI, Nature **614**, 676 (2023)] for the repetition code, giving competitive decoding accuracies. Although training is computationally demanding, inference is fast and scales approximately linearly with the space-time volume of the code, making this a potentially competitive approach to future integrated systems.

Extreme photonic sensing

Miroslav Ježek

Palacky University Olomouc

We report on advanced detectors and sensors of statistical, vector, and correlation properties of light down to the single-photon level. The novel detection schemes harness complex optical networks and deep learning. We present minimum-error multi-projection quantum measurement, resource-efficient polarization sensor and entanglement detector, and single-emitter superresolution sensing. We discuss the efficiency and measurement-device independence of the presented detection modalities.

Improving NISQ via quantum middleware

Maciej Koch-Janusz

Haiqu, Inc. / University of Zurich

A useful quantum computing system is impossible without tailored software, the careful design of which is critical especially in the NISQ era. In this short talk I will introduce Haiqu, a quantum middleware startup focused on improving execution of quantum circuits on current noisy devices, and showcase one of our open-source tools, the Rivet transpiler, and its functionalities which may be useful in QML applications.

The first 10 years of AI-designed quantum experiments

Mario Krenn

Max Planck Institute for the Science of Light

On a cold night in March 2014, the file “solution.txt” suddenly appeared on a computer. This file contained the answer to a question that human scientists had not yet been able to answer: the blueprint of an experimental design to create a novel, complexly entangled quantum system. What the computer discovered ultimately led to the first observation of new complex quantum correlations in the laboratory.

In my talk, I trace the developments in this dynamic research field, from advanced applications such as quantum optics, innovative microscopes and gravitational wave detectors as well as some of the crucial underlying technologies, involving auto-differentiation and large-scale exploration.

Associative memories in the Quantum Regime

Adrià Labay Mora

IFISC

Associative memories are brain-inspired computational systems that can solve and model a wide range of tasks, ranging from pattern and speech recognition to big data analysis. A system equipped with associative memory can identify the stored pattern that is most similar to the clue according to a properly defined distance. We study a nonlinear dissipative driven-dissipative oscillator able to perform pattern discrimination in a metastable phase where patterns are encoded in coherent or squeezed states. We show that coherent states can exceed the limiting storage capacity of Hopfield networks in both classical and quantum regimes. Additionally, we introduce the possibility of storing patterns as purely quantum states, with no classical analogue. Finally, we generalise the framework of associative memories by investigating the shape of quantum channels that perform such tasks.

Artificial discovery of coupled-mode circuits with desired scattering behavior

Jonas Landgraf

Max Planck Institute for the Science of Light

In quantum optics, a small number of building blocks, like resonators, waveguides, driving-induced couplings, and parametric interactions allow the design of a broad variety of devices and functionalities, distinguished by their scattering properties. These include transducers, amplifiers, and nonreciprocal devices, like isolators or circulators.

Usually, the design of such a system is handcrafted by an experienced scientist. In our work, we develop a discovery algorithm that automates this process. By optimizing the continuous and discrete system properties our automated search identifies the minimal resources required to realize the requested scattering behavior. The discovered architectures represent classes of solutions and are not bound to special numerical values or platforms. Our approach is applicable to optical, microwave, mechanical, electrical, and hybrid circuits.

Better quantum computing via reinforcement learning

Florian Marquardt

Max Planck Institute for the Science of Light

In this talk I will give a brief overview of some of our recent insights in exploiting reinforcement learning for improving quantum computation. These include the optimization of ZX diagrams, the discovery of fault-tolerant quantum state preparation, and better feedback strategies for GKP code quantum error correction.

Hybrid quantum neural networks: from theory to applications

Alexey Melnikov

Terra Quantum

A parametrized quantum circuit represents a fundamental unit of a learnable quantum information processing system, characterized by its capability to adaptively modify a state of the quantum system in response to empirical data. By integrating quantum circuits with conventional neural network architectures, hybrid quantum neural networks represent an approach that leverages the capabilities of quantum systems to enhance machine learning solutions. This talk discusses the capabilities and limitations of the hybrid quantum neural networks. First, different ways to hybridize quantum and classical neural networks will be shown. Second, this hybridization will be analyzed in terms of its capacity and limitations. Third, the choice of hardware backends and SDKs will be discussed. Fourth, applications in industrial tasks will be presented. This synergy in hybrid quantum neural networks pushes the boundaries of what is potentially achievable with near-term quantum computers.

Diffusion models: from quantum circuit synthesis to experiment generation

Gorka Muñoz-Gil

University of Innsbruck

Quantum computing has recently emerged as a transformative technology. Yet, its promised advantages rely on efficiently translating quantum operations into viable physical realizations. In this work, we use generative machine learning models, specifically denoising diffusion models (DMs), to facilitate this transformation. Leveraging text-conditioning, we steer the model to produce desired quantum operations within gate-based quantum circuits. Notably, DMs allow to sidestep during training the exponential overhead inherent in the classical simulation of quantum dynamics—a consistent bottleneck in preceding ML techniques. We demonstrate the model's capabilities across two tasks: entanglement generation and unitary compilation. The model excels at generating new circuits and supports typical DM extensions such as masking and editing to, for instance, align the circuit generation to the constraints of the targeted quantum device. Given their flexibility and generalization abilities, we envision DMs as pivotal in quantum circuit synthesis, enhancing both practical applications but also insights into theoretical quantum computation.

Unifying (Quantum) Statistical and Parametrized (Quantum) Algorithms

Alexander Nietner

Freie Universität Berlin

Kearns' statistical query (SQ) oracle (STOC'93) lends a unifying perspective with a rich explanatory power for most classical machine learning algorithms. This ceases to be true in quantum learning, where many settings admit neither an SQ nor a quantum statistical query (QSQ) analog. In this work, we take inspiration from Kearns' SQ oracle and Valiant's weak evaluation oracle (TOCT'14) to unify statistical and parametrized learning. We present a systematic study of the problem of learning from an evaluation oracle, which provides an estimate of function values, and extend Feldman's framework for SQ learning (COLT'17) to our setting. This leads to unconditional lower bounds and a characterization of learning linear function classes which are directly applicable to the QSQ setting and virtually any algorithm based on loss function optimization.

This formalism allows to extend prior results on the learnability of quantum states and their output distributions from the SQ to the (multi-copy) QSQ setting, implying exponential separations between learning stabilizer states from (multi-copy) QSQ's versus from quantum samples.

This talk will focus on the application of our formalism to better understand the hardness of quantum machine learning (QML). We will share a broad understanding of the hardness of various QML tasks, which goes beyond the intuition of barren plateaus (BP) and enables us to separate the implications of BP's depending on the context in which they appear.

Discovering quantum error correction strategies with reinforcement learning

Jan Olle

Max Planck Institute for the Science of Light

Finding optimal ways to protect quantum states from noise remains an outstanding challenge across all quantum technologies, and quantum error correction (QEC) is the most promising strategy to address this issue. Constructing QEC codes is a complex task that has historically been powered by human creativity with the discovery of a large zoo of families of codes. However, in the context of real-world scenarios there are two challenges: these codes have typically been categorized only for their performance under an idealized noise model and the implementation-specific optimal encoding circuit is not known. In this work, we train a deep reinforcement learning agent that automatically discovers both QEC codes and their encoding circuits for a given gate set, qubit connectivity, and error model. I will introduce and discuss the concept of a noise-aware meta-agent, which learns to produce encoding strategies simultaneously for a range of noise models, thus leveraging transfer of insights between different situations. Moreover, thanks to the use of the stabilizer formalism and a vectorized Clifford simulator, our RL implementation is extremely efficient, allowing us to produce many codes and their encoding circuits from scratch within seconds, with code distances varying from 3 to 5 and with up to 20 physical qubits. Finally, I will discuss how this strategy can in principle be scaled up to 40 to 100 qubits and code distances from 6 to 10.

What algorithms should we study with 100 qubits and 1M logical gates?

Alexandru Paler

Aalto University

The recent neutral atom error-correction experiments have sparked the interest in the implementation and execution of larger scale protocols and computations. The near-term research question is reflected in the title of this talk, and in order to answer it, the most feasible approach is co-design. This means that QEC and FT-protocols, on the one hand, and algorithms/circuits, on the other hand, influence each other through the architecture of the hardware. Co-design is a complex process and can be both theoretically and practically, investigated by analyzing the software stack that translates an algorithm to an executable circuit. The stack includes levels which are not necessarily traversed in a top-down, sequential manner [9]: a) compilation into a fault-tolerant form [1,2], b) preparation for error-correction [3,4], and c) various types of optimization [5,6]. Computational faults are avoided and errors are corrected during the execution of the circuit [7].

In order to co-design algorithms with hundreds of qubits and millions of gates, one should start from the following research questions related to the execution of simpler protocols: a) how are injection protocols (e.g. [8]) reflected in the decoding of correlated errors?; b) do logical qubits suffer from novel/unexpected types of errors, and if so, what is the effect of these errors on the structure of the fault-tolerance compilation primitives? c) what logical cycle times are to be expected based on the underlying architecture, and how much of an improvement is necessary for lowering the resource counts?

In the first part of the talk, we will present a realistic software stack that is already available and can be used for automating the search and co-design of algorithms. In order to answer the research questions, new models, methods and (software) tools need to be researched and implemented. These aspects are discussed in the second part of the talk. The methods and results presented herein were developed partially within projects (where the speaker is a PI) funded by the DARPA Quantum Benchmarking program, QuantERA and Google.

References

1. Vijayan, M.K., Paler, A., Gavriel, J., Myers, C.R., Rohde, P. and Devitt, S.J., 2024. Compilation of algorithm-specific graph states for quantum circuits. *Quantum Science and Technology*.
2. Paler, A., Polian, I., Nemoto, K. and Devitt, S.J., 2017. Fault-tolerant, high-level quantum circuits: form, compilation and description. *Quantum Science and Technology*, 2(2), p.025003.

3. Watkins, G., Nguyen, H.M., Seshadri, V., Watkins, K., Pearce, S., Lau, H.K. and Paler, A., 2023. A High Performance Compiler for Very Large Scale Surface Code Computations. arXiv preprint arXiv:2302.02459.
4. Paler, A., Fowler, A.G. and Wille, R., 2017. Synthesis of arbitrary quantum circuits to topological assembly: Systematic, online and compact. Scientific reports, 7(1), p.10414.
5. Moflic, I. and Paler, A., 2023, December. Towards Faster Reinforcement Learning of Quantum Circuit Optimisation: Exponential Reward Functions. In Proceedings of the 18th ACM International Symposium on Nanoscale Architectures (pp. 1-2).
6. Paler, A., 2022. Quantum Fourier addition simplified to Toffoli addition. Physical Review A, 106(4), p.042444.
7. Paler, A. and Fowler, A.G., 2023. Pipelined correlated minimum weight perfect matching of the surface code. Quantum, 7, p.1205.
8. Gavriel, J., Herr, D., Shaw, A., Bremner, M.J., Paler, A. and Devitt, S.J., 2023. Transversal injection for direct encoding of ancilla states for non-Clifford gates using stabilizer codes. Physical Review Research, 5(3), p.033019.
9. Paler, A., 2020. Aggregated Control of Quantum Computations: When Stacked Architectures Are Too Good to Be Practical Soon. Computer, 53(8), pp.74-78..

Quantum variational re-uploading for machine learning

Adrián Pérez Salinas

Leiden University

Variational quantum algorithms have emerged as a promising framework for Noisy Intermediate-Scale Quantum computing. While originally conceived to find ground states, VQAs were modified to include data through embedding maps, opening the path to variational quantum machine learning. The re-uploading scheme provided a recipe to systematically introduce data, allowing for optimization of the embeddings. In this talk we will review the re-uploading framework, with in-depth focus on the expressivity and trainability of the functions output by the machine learning models. We also explore the connection to dataless variational models. Finally, we will discuss about educated guesses for future research

Measurement-based Quantum Computation for Machine Learning

Hendrik Poulsen Nautrup

University of Innsbruck

In this talk we explore measurement-based quantum computation (MBQC) for machine learning. To this end, we show that unitary MBQC on a regular lattice with measurements in the XY-basis is equivalent to a circuit model quantum computation based on Clifford quantum cellular automata. While this model of quantum computation is ideally suited as a Hardware efficient Ansatz for certain quantum computation architectures such as neutral atoms, there remains sufficient flexibility to use it as a problem specific Ansatz for variational quantum circuits. We further explore non-unitary MBQC as a variational Ansatz for generative modelling and find that it outperforms the corresponding unitary Ansätze both numerically and algebraically.

Real data classification problems with a few qubits

Pablo Rodriguez

University of the Basque Country (UPV/EHU)

Kernel methods are pivotal in machine learning, with Embedding Quantum Kernels (EQKs) showing great promise in quantum systems. However, selecting the optimal embedding for EQKs poses a challenge. Our solution involves a p -qubit Quantum Neural Network (QNN) utilizing data re-uploading to determine the optimal q -qubit EQK for a given task (p -to- q). This approach constructs the kernel matrix just once, enhancing efficiency. We concentrate on n -to- n and 1 -to- n , introducing a scalable method for training an n -qubit QNN and demonstrating the utility of a single-qubit QNN in constructing potent EQKs. We demonstrate the remarkable capability of very small quantum systems (2,3 qubits) to construct complex quantum machine learning models, showcasing their effectiveness in performing classification tasks, even on real-world data.

Classical Surrogate Simulation of Quantum Systems using LOWESA

Manuel S. Rudolph

EPFL

We introduce a new classical algorithm - LOWESA - for simulating quantum systems. This algorithm constructs an analytical surrogate landscape for operator expectation values of quantum circuits that consists of weighted sums of trigonometric functions ordered from low to high-frequency. While constructing the surrogate landscape does incur an initial overhead, one feature of LOWESA is that re-evaluating it at different rotation parameters is exceptionally fast compared to other classical simulation methods. One could, for example, rapidly study the time evolution of entire families of Hamiltonians, initial states, and target observables for varying simulation time, due to these properties translating to different parameters in the corresponding time evolution quantum circuits. Fast expectation and gradient evaluations additionally enable variational approaches or circuit compilation, which can readily be utilized and improved by other classical algorithms with different trade-offs or quantum algorithms once the hardware permits it.

As a case study, we simulate the time evolution of 127 qubits in a transverse-field Ising model on a heavy-hexagon lattice with up to 20 Trotter steps (see Nature 618, 500-505 (2023)). Our approach vastly outperforms matrix product state simulation, and is generally on-par with other leading methods in terms of precision, but yields significantly faster execution time per evaluation. Beyond the promising empirical results, we provide initial error bounds for LOWESA in the cases of noisy and noise-free simulation.

Time series forecasting using quantum machine learning

Asel Sagingalieva

Terra Quantum

The use of quantum neural networks for machine learning is a paradigm that has recently attracted considerable interest. Under certain conditions, these models approximate the distributions of their datasets using truncated Fourier series. Owing to the trigonometric nature of this fit, angle-embedded quantum neural networks have an advantage in fitting harmonic features in a given dataset. This talk discussed Hybrid Quantum Long Short-Term Memory and Hybrid Quantum Sequence-to-Sequence neural network architectures that are applied in industrial time series forecasting settings.

Machine learning for trapped ion quantum information experiments

Philipp Schindler

University of Innsbruck

TBD.

Fully autonomous tuning of a spin qubit

Jonas Schuff

University of Oxford

Spanning over two decades, the study of qubits in semiconductors for quantum computing has yielded significant breakthroughs [1–3]. However, the development of large-scale semiconductor quantum circuits is still limited by challenges in efficiently tuning and operating these circuits. Identifying optimal operating conditions for these qubits is complex, involving the exploration of vast parameter spaces [4]. This presents a real ‘needle in the haystack’ problem, which, until now, has resisted complete automation due to device variability and fabrication imperfections [5]. In this study, we present the first fully autonomous tuning of a semiconductor qubit, from a grounded device to Rabi oscillations, a clear indication of successful qubit operation. We demonstrate this automation, achieved without human intervention, in a Ge/Si core/shell nanowire device. Our approach integrates deep learning, Bayesian optimization, and computer vision techniques. We expect this automation algorithm to apply to a wide range of semiconductor qubit devices, allowing for statistical studies of qubit quality metrics. As a demonstration of the potential of full automation, we characterise how the Rabi frequency and g-factor depend on barrier gate voltages for one of the qubits found by the algorithm. Twenty years after the initial demonstrations of spin qubit operation, this significant advancement is poised to finally catalyze the operation of large, previously unexplored quantum circuits.

References

- [1] D. Loss and D. P. DiVincenzo, *Phys. Rev. A*, vol. 57, no. 1, p. 120 (1998)
- [2] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard, *Science*, vol. 309, no. 5744, pp. 2180–2184 (2005)
- [3] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. Vandersypen, *Nature*, vol. 601, no. 7893, pp. 343–347 (2022)
- [4] S. Neyens, O. Zietz, T. Watson, F. Luthi, A. Nethwewala, H. George, E. Henry, A. Wagner, M. Islam, R. Pillarisetty, et al., *arXiv preprint arXiv:2307.04812* (2023)
- [5] H. Moon, D. Lennon, J. Kirkpatrick, N. van Esbroeck, L. Camenzind, L. Yu, F. Vigneau, D. Zumbühl, G. A. D. Briggs, M. Osborne, et al., *Nat. Commun.*, vol. 11, no. 1, pp. 1–10 (2020).

Enhancing Privacy and Security in the Quantum World

Andreas Trügler

Know-Center, Graz University of Technology

Quantum computing stands at the cusp of revolutionizing information technology by offering computational opportunities far exceeding classical computers. The growing and seamless availability of quantum computing resources through cloud infrastructures marks a significant leap forward regarding user-friendliness and accessibility of quantum hardware. However, this trend also heralds novel privacy concerns. As individuals and organizations begin to access quantum cloud capabilities, they inevitably find themselves entrusting sensitive data and quantum algorithms to third-party providers.

In this presentation we delve into the intricacies of privacy and security in the realm of quantum computation, focusing on two pivotal aspects: Side-channel attacks on quantum hardware and the possibility to merge concepts from classical cryptography with quantum computation. We present recent findings on how quantum decoherence processes and information leakage from soft- and hardware can be exploited to reconstruct sensitive data or details of the computed quantum algorithm. Additionally, we explore concepts from classical cryptography and discuss how integration of quantum computation protocols can lead to novel approaches in the field that open new door towards the next generation of secure quantum powered applications for complex computational tasks.

In conclusion, this presentation aims to stimulate further research and collaboration and to shed light on the critical challenges and emerging solutions at the intersection of privacy, security, and quantum computation, especially with respect to machine learning applications.

Investigating the optimization landscape of neural-network quantum states

Agnes Valenti

Flatiron Institute, CCQ

Neural-network wave-functions have been proven to serve as flexible variational ansatz for the ground state search of many-body Hamiltonians. At the same time, they are notoriously hard to optimize, in particular in the presence of frustration. The interplay between trainability and expressivity of the ansatz is further affected by computational limitations that pose constraints on the number of parameters and/or the number of samples. We investigate this broad question in a systematic approach, by comparing the optimization landscape of some of the most common neural-network architectures (RBMs, RNNs and GCNNs). In particular, we analyze the influence of sampling and parametrization on the training dynamics and expressivity of the respective architectures, on a use-case of a frustrated Hamiltonian: the J1-J2 model on a triangular lattice.

All-fiber microendoscopic polarization sensing at single-photon level aided by deep-learning

Dominik Vasinka

Palacky University Olomouc

The polarization of light conveys crucial information about the spatial ordering and optical properties of a specimen. However, precise polarization measurement under challenging conditions, including constrained spaces, low light levels, and high-speed scenarios, represents a severe challenge. Addressing this problem, we introduce a real-time polarization measurement method accurate down to a single-photon level that provides complete information about the polarization state. Free of moving components, the polarization sensor utilizes a few-mode fiber followed by a fiber array and a detector array. The calibration of the sensor relies on a deep learning model yielding unprecedented accuracy across all polarization states, including partially polarized light. We validate the approach by visualizing the polarization structure of a biological specimen. Our method offers an efficient and reliable solution for real-time polarization sensing and microendoscopy under low-light conditions.

Advances in quantum simulation with Neural Quantum States

Filippo Vicentini

École Polytechnique

Machine Learning algorithms have advanced tremendously in the last two decades, tackling difficult problems such as image and natural language recognition, classification or generative synthesis. Part of the success owes to the development of variational approximators which can efficiently approximate unknown functions living in high-dimensional spaces: Neural Networks.

By today, many ground-state problems can be solved with extremely high accuracy using such techniques, but simulating dynamical properties still presents several challenges. For example last summer we have highlighted that established algorithms based on time-dependent variational principles are biased and cannot be used reliably to simulate dynamics, while unbiased alternatives are computationally expensive.

In this seminar I will discuss recent advances and open challenges within the domain in the case of both closed and open quantum systems.

Learning quantum processes with quantum statistical queries

Chirag Wadhwa

University of Edinburgh

In this talk, we will introduce a formal framework for learning quantum processes from statistical queries (QPSQs). The framework allows us to propose an efficient QPSQ learner for the outputs of arbitrary quantum processes accompanied by a provable performance guarantee. We also provide numerical simulations to demonstrate the efficacy of this algorithm. In our new framework, we prove exponential query complexity lower bounds for learning unitary 2-designs, and a doubly exponential lower bound for learning haar-random unitaries. The practical relevance of this framework is exemplified through application in cryptography, where we apply our learning algorithms and lower bounds to characterize the security of Classical Readout Quantum Physical Unclonable Functions (CR-QPUFs), addressing an important open question in the field of quantum hardware security. The talk will be based on [arXiv:2310.02075v2](https://arxiv.org/abs/2310.02075v2).

Fisher Information and Artificial Neural Networks

Maximilian Weimar

TU Wien

Artificial neural networks (ANNs) are powerful tools for learning complicated, non-linear connections between measurements and targeted observables from data. For instance, in physical experiments, they are widely used to reconstruct properties of a system from the measured data, which is inherently noisy due to technical limitations or fundamental quantum effects, such as fluctuations in the electromagnetic field. This noise ultimately limits the performance of an ANN as it learns properties of the system. An important question to ask is whether this limit can be identified and whether the network can perform at it.

For a broad class of physical problems that entail the estimation of a continuous parameter of a distribution from a data set, the so-called Fisher information (FI) [1] determines this fundamental bound. Although this was known for a long time, so far, the FI was only accessible for data following very simple distributions. To deal with the most general cases, where the distribution is complicated and unknown, we introduce a data-driven method [2] for approximating the FI. We collect data from a scattering experiment that is dominated by uncontrolled noise and train ANNs to predict a physical parameter from the data. We then compare their performances with the bound we derived and demonstrate that the physical parameter estimation task can be performed optimally [3].

In a next step, we extend our analysis to the hidden layers of the ANN. Not only does our method allow us to access the FI of the input data but it also applies to the complicated, high-dimensional representations of the data that are generated deep within the network. We interpret the FI as a quantity that flows through the ANN from layer to layer, while improper weights and biases manifest as sinks for the FI flow. Observing the FI during the training us to peek into the black box [4] of the ANN.

References:

- [1] A. Van den Bos, Parameter estimation for scientists and engineers (John Wiley & Sons, 2007).
- [2] M. Weimar, L. Rachbauer, S. Rotter, Fisher information flow in artificial neural networks (to be published)
- [3] I. Starshynov, et al., Retrieving information from random physical systems: the ultimate limit of deep neural networks (to be published)
- [4] N. Tishby and N. Zaslavsky, Deep learning and the information bottleneck principle, in 2015 IEEE information theory workshop (ITW) (IEEE, 2015) pp. 1–5.

Tight and Efficient Gradient Bounds for Parameterized Quantum Circuits

Christa Zoufal

IBM Quantum

The training of a parameterized model largely depends on the landscape of the underlying loss function. In particular, vanishing gradients (also known as barren plateaus) are a central bottleneck in the scalability of variational quantum algorithms (VQAs), and are known to arise in various ways, from circuit depth and hardware noise to global observables. However, a caveat of most existing gradient bound results is the requirement of t -design circuit assumptions that are typically not satisfied in practice. In this work, we loosen these assumptions altogether and derive tight upper and lower bounds on loss and gradient concentration for a large class of parameterized quantum circuits and arbitrary observables, which are significantly stronger than prior work. Moreover, we show that these bounds, as well as the variance of the loss itself, can be estimated efficiently and classically – providing practical tools to study the loss landscapes of VQA models, including verifying whether or not a circuit/observable induces barren plateaus. In particular, our results can readily be leveraged to rule out barren plateaus for a realistic class of ansätze and mixed observables, namely, observables containing a non-vanishing local term. This insight has direct implications for hybrid Quantum Generative Adversarial Networks (qGANs), a generative model that can be reformulated as a VQA with an observable composed of local and global terms. We prove that designing the discriminator appropriately leads to 1-local weights that stay constant in the number of qubits, regardless of discriminator depth. Combined with our first contribution, this implies that qGANs with shallow generators can be trained at scale without suffering from barren plateaus – making them a promising candidate for applications in generative quantum machine learning. We demonstrate this result by training a qGAN to learn a 2D mixture of Gaussian distributions with up to 16 qubits, and provide numerical evidence that global contributions to the gradient, while initially exponentially small, may kick in substantially over the course of training.

Posters

Analogue Quantum Simulators - table - top experiments for quantum field theory

Nataliia Bazhan

Atominstitut, TU Wien

Analogue (quantum) simulators have gained increasing traction over the last decade as an interdisciplinary branch between experimental and theoretical physics. They present an ideal platform to study fundamental questions of quantum field theory, which are otherwise not or only indirectly accessible in experiments. We will present general ideas on analogue gravity and cosmology based on classical and quantum fluids, as well as our recent advances on analogue QFT simulators in and beyond the linear regime of the sine-Gordon model, in particular in connection with optimal control theory for continuous field theories and the opportunities and challenges towards high-precision experiments of quantum field theory in table-top experiments.

Constrained and Vanishing Expressivity of Quantum Fourier Models

Mario Herrero Gonzalez

University of Edinburgh

In this work, we highlight an unforeseen behavior of the expressivity of Parameterized Quantum Circuits (PQC) for machine learning. A large class of these models, seen as Fourier Series which frequencies are derived from the encoding gates, were thought to have their Fourier coefficients determined by the trainable gates. Here, we demonstrate a new correlation between the Fourier coefficients of the quantum model and its encoding gates via the property of `\textit{frequency redundancies}`. In addition, we display a phenomenon of vanishing expressivity in certain settings, where the variance of some Fourier coefficients vanish exponentially when the number of qubits grows. These two behaviors imply novel forms of constraints which limit the expressivity of PQCs, and therefore imply a new inductive bias for quantum models.

Qubit state tomography with incomplete measurements using neural networks

Francis Marcellino

University of Geneva

Qubit state tomography involves inferring the density matrix of a quantum state from a set of measurements in different bases. The size of the set of measurements required for full tomography is exponential in the number of qubits, but reasonably accurate reconstructions can nevertheless be performed using many fewer measurements, especially for states of low rank. We explore the use of neural networks with various architectures for this task, and show that they can offer an advantage over traditional tomography techniques.

Compiling near-term fault tolerant quantum experiments

Ioana Moflic

Aalto University

We are reaching the starting point of the post-NISQ era where there are enough physical qubits for executing the first fault-tolerant experiments using logical error-corrected qubits. To this end, automatic preparation and evaluation of the near-term experiments is becoming an important tool.

I will talk about software methods for optimizing Clifford+T circuits. My work focused on novel and scalable compilation techniques:

- reinforcement learning (RL) for quantum circuit optimisation [1,2];
- graph neural network autoencoders for compressing the quantum circuit representations [3];
- an extremely scalable, multi-threaded software tool used to analyze and optimize quantum circuits with hundreds of millions of gates [4].

The latter is still under development and one of its applications is to automatically compile and evaluate the fault-tolerance of circuits protected by flag qubits.

References

- [1] **Moflic, I.** and Paler, A., 2023, December. Towards Faster Reinforcement Learning of Quantum Circuit Optimisation: Exponential Reward Functions. In Proceedings of the 18th ACM International Symposium on Nanoscale Architectures (pp. 1-2).
- [2] **Moflic, I.** and Paler, A., 2023, December. Cost Explosion for Efficient Reinforcement Learning Optimisation of Quantum Circuits. In 2023 IEEE International Conference on Rebooting Computing (ICRC) (pp. 1-5). IEEE.
- [3] **Moflic, I.**, Garg, V. and Paler, A., 2023. Graph Neural Network Autoencoders for Efficient Quantum Circuit Optimisation. arXiv preprint arXiv:2303.03280.
- [4] **Moflic, I.** and Paler, A., 2024. Optimising Circuits with Millions of Gates Within a Few Seconds, manuscript submitted at QCTIP 2024, Edinburgh.

Thermometry of one-dimensional Bose gases with neural networks

Frederik Møller

Atominstitut, TU Wien

One-dimensional Bose gases trapped on an atom chip represent a quantum analog simulator of several important quantum field theory models. To probe the experimental system, absorption images of the atomic cloud following free expansion are taken. Averaging over many images, expectation values and correlations functions can be extracted and compared to known theories. By virtue of the complicated nature of such systems the recorded absorption images contains a large amount of information, much of which is discarded in the extraction process. However, through machine learning techniques more efficient analysis of the experimental data may be possible.

As a demonstration, we train a neural network to predict the temperature of a one-dimensional Bose gas given a single absorption image. We compare our model to the established method of density ripple thermometry, which requires several tens of images to produce accurate results. Here we find the predictions of the two methods compatible, with the neural network reaching similar precision by combining the predictions of much fewer images.

Real data classification problems with a few qubits

Pablo Rodriguez

University of the Basque Country (UPV/EHU)

Kernel methods are pivotal in machine learning, with Embedding Quantum Kernels (EQKs) showing great promise in quantum systems. However, selecting the optimal embedding for EQKs poses a challenge. Our solution involves a p -qubit Quantum Neural Network (QNN) utilizing data re-uploading to determine the optimal q -qubit EQK for a given task (p -to- q). This approach constructs the kernel matrix just once, enhancing efficiency. We concentrate on n -to- n and 1 -to- n , introducing a scalable method for training an n -qubit QNN and demonstrating the utility of a single-qubit QNN in constructing potent EQKs. We demonstrate the remarkable capability of very small quantum systems (2,3 qubits) to construct complex quantum machine learning models, showcasing their effectiveness in performing classification tasks, even on real-world data.

Machine Learning for Optimal Quantum Solids Deposition

Aslihan Sasmaz

Atominstitut, TU Wien

Quantum solids made of noble gases (Ar, Ne, He, etc.) or molecules (p-H₂, N₂, etc.) provide a platform for studying complex systems such as quantum many-body systems. These cryocrystals are weakly interacting and magnetically quiet, which makes them a suitable host for impurities. We deposit quantum solids of neon (Ne) with sodium (Na) impurities inside the cryostat directly atop superconducting circuit, to achieve strong coupling of alkali atoms and a superconducting resonator.

However, the final properties of each cryocrystal depend on five growing parameters. We aim to find the best parameters using machine learning algorithms. I will present the current progress on the optimization of quantum solids deposition.

Universal scaling dynamics in Bose gases far from equilibrium

Rugway Wu

Atominstitut, TU Wien

We present experimental results of universal scaling behaviour during the evolution and relaxation of isolated quantum many-body systems quenched far from equilibrium. Different experimental scenarios are explored:

- [Rubidium: Quasi-1D] Defect nucleation and its connection to the Kibble-Zurek mechanism for rapidly cooling quenched one-dimensional Bose gases and coupled quantum wires.
- [Metastable Helium: Quasi-1D] Differences in the self-similar scaling behaviour between the condensed and thermal parts of the gas during condensation.
- [Lithium Molecules: Elongated 3D] The appearance of multiple self-similar scaling regimes for strongly interacting gases. Tuneable via the atomic interaction strength, the time dependence of scaling exponents might be connected to prescaling and/or the approach of multiple non-thermal fixed points during the relaxation of strongly interacting systems.

Quantum Kernel Methods on the photonic system

Zhenghao Yin

University of Vienna

In recent years, machine learning has had a remarkable impact on standard computation, with applications ranging from scientific to everyday-life scopes. However, as the complexity of the addressed task grows, energy consumption and computational power requirements become a bottleneck.

In this context, it has been shown that quantum (or quantum-inspired) computation enables for higher efficiency.

Here, we demonstrate a kernel method on a photonic integrated processor to perform a binary classification task, by exploiting quantum interference.

We benchmark our protocol against standard algorithms and show that it outperforms them for the given tasks.

Even in a regime that does not display quantum interference and exploits only the coherence of single photons, we achieve a significant improvement with respect to standard algorithms.

Both versions of our protocol allow the implementation of effective classifications while only utilizing linear-optical elements.

This means that we do not need entangling gates and can readily raise the dimension of our system through additional circuit modes and/or injected photons.

Our result opens the way to more efficient computing algorithms and for the formulation of tasks where quantum effects enhance the effectiveness of standard methods.

Searching Optimal Trajectories for Condensate Splitting

Tiantian Zhang

Atominsitut, TU Wien

Bose-Einstein condensates (BECs) in a double well function as an excellent matter-wave interferometer with enhanced phase sensitivity originating from the nonlinear inter-atom interactions [1]. By optimising the BEC splitting procedure, we can prepare BECs in a double well more time efficiently and without introducing unwanted motional excitations. We generate optimised trajectories using the off-line iterative learning control (ILC) method [2] based on a one-dimensional mean-field simulation. We have observed suppression of induced breathing and sloshing from a fast ramp. Furthermore, we have investigated experimentally how the quantum properties, such as number squeezing, evolve [1] after such optimised ramps. Intrigued by these observations, we envision an extended ability to engineer arbitrary many-body quantum states of BECs in a double well.

[1] Zhang, T., Maiwöger, M., Borselli, F., Kuriatnikov, Y., Schmiedmayer, J., & Prüfer, M. (2023).

Squeezing oscillations in a multimode bosonic Josephson junction. ArXiv. /abs/2304.02790

[2] Deutschmann-Olek, A., Tajik, M., Calzavara, M., Schmiedmayer, J., Calarco, T., & Kugi, A. (2022,

December). Iterative shaping of optical potentials for one-dimensional Bose-Einstein condensates. In

2022 IEEE 61st Conference on Decision and Control (CDC) (pp. 5801-5806). IEEE.