Quantum Physics and Machine Learning

Universitätszentrum Obergurgl

13th - 19th April 2025

QUANDELA

Meta-Designing Quantum Experiments with Language Models

Sören Arlt

Max Planck Institute for the Science of Light

Artificial Intelligence (AI) can solve complex scientific problems beyond human capabilities, but the resulting solutions offer little insight into the underlying physical principles. One prominent example is quantum physics, where computers can discover experiments for the generation of specific quantum states, but it is unclear how finding general design concepts can be automated. In this work we address this challenge by training a transformer-based language model to create human-readable Python code, which solves an entire class of problems in a single pass. This strategy, which we call meta-design, enables scientists to gain a deeper understanding and extrapolate to larger experiments without additional optimization.

Simulations of Malicious Cross-talk for Side-channel Exploitation in Parallel Quantum Computing

Brennan Bell

Know-Center GmbH

Cross-talk is the phenomenon by which qubit operations cause changes in behaviour to one another during a circuit execution. These changes can occur during qubit initialization, gate operations, and measurement readouts, but our focus lies on gate-oriented cross-talk operations which can modeled by the Kraus operators. In this talk, we will discuss a series of experiments using offline Qiskit simulators, concentrating on both hand-crafted and device-profiled noise models. We investigate how parallel execution on a single device by two parties can lead to information leakage via cross-talk, and aim to use this as a starting point for a discussion regarding instigation (malicious) and mitigation strategies.

Ignorance, Thermodynamics, and Learning

Andrew Briggs

University of Oxford

In a conventional deep neural network for AI, the values of the weights are scarcely more knowable than the position and momentum of every molecule in a container of gas. For Avogadro scale systems, ignorance of individual particles is complete and statistical knowledge determines the physical observables and efficiencies of machines. In sufficiently small systems it may be possible to know the position and momentum of every relevant particle. In mesoscopic systems fluctuations become significant. Szilard engines can be built, Landauer erasure can be measured, and the thermodynamic cost of a clock can be determined. Probability remains subjective.

This prompts the question how thermodynamics and statistical mechanics can elucidate processes in AI. For example, in a nanoscale learning machine, such as might be constructed from one or a few qubits, it may be possible to know the parameters of individual activation switches and hence study their thermodynamic behaviour. It then becomes possible to apply fluctuation theorems, and hence elucidate the entropic cost of the learning process. This applies to both classical and quantum learning machines, and may provide new insights into the way how the ‘knowledge’ embodied in the weights of a network is ‘learned’ and how the size of a system may affect how different levels of subjective ignorance can still generate reliable outputs.

Efficient Online Quantum Circuit Learning with No Upfront Training

Piotr Czarnik

Jagiellonian University

We propose a surrogate-based method for optimizing parameterized quantum circuits which is designed to operate with few calls to a quantum computer. We employ a computationally inexpensive classical surrogate to approximate the cost function of a variational quantum algorithm. An initial surrogate is fit to data obtained by sparse sampling of the true cost function using noisy quantum computers. The surrogate is iteratively refined by querying the true cost at the surrogate optima, then using radial basis function interpolation with existing and new true cost data. The use of radial basis function interpolation enables surrogate construction without hyperparameters to pre-train. Additionally, using the surrogate as an acquisition function focuses hardware queries in the vicinity of the true optima. For 16-qubit random 3-regular Max-Cut problems solved using the QAOA ansatz, we find that our method outperforms the prior state of the art. Furthermore, we demonstrate successful optimization of QAOA circuits for 127-qubit random Ising models on an IBM quantum processor using measurement counts of the order of 10^4-10^5. The strong empirical performance of this approach is an important step towards the large-scale practical application of variational quantum algorithms and a clear demonstration of the effectiveness of classical-surrogate-based learning approaches.

Quantum memory advantages for stochastic simulation

Thomas Elliott

University of Manchester

Simulating quantum dynamics on a classical computer bears a resource cost that grows exponentially with the size of the system, and even the simplest of quantum systems often exhibit seemingly complex behaviors. This apparent problem can be recast as a positive - complex classical systems can be simulated efficiently on simple quantum computers. In this talk I will discuss the application of quantum technologies to the modelling of stochastic processes, for which quantum simulators can operate with lower memory cost than any classical alternative, in both lossless and lossy compression settings. Particularly, I will highlight examples of quantitative scaling divergences in modelling highly non-Markovian processes, wherein the provably-memory-minimal classical simulator must store diverging amounts of information with increasing precision, while arbitrary precision can be achieved with a finite-sized quantum simulator. I will also discuss recent works on the experimental implementation of such quantum memory advantages, and the extension to modelling adaptive agents.

Gradient-descent methods for fast quantum state tomography

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 my talk last year, I presented results on applying generative adversarial neural networks to quantum state tomography [1,2,3], and on constrained gradient-descent optimization for quantum process tomography [4,5]. While both approaches were successful, outperforming standard methods in terms of both the amount of time and data needed, the question was raised whether constrained gradient-descent optimization could be applied to quantum state tomography and do even better than the machine-learning algorithm there. In this talk, I answer that question in the affirmative. I will present recent results [6] showing how gradient descent can be applied to quantum state tomography using several parameterisations that ensure the estimated state is a physically valid density matrix. I will show that these parameterisations enable control of the rank of the ansatz, and that these algorithms effectively handle noisy and incomplete data sets, yielding significantly higher reconstruction fidelity than state-of-the-art methods that we benchmark against.

[1] S. Ahmed et al., Phys. Rev. Lett. 127, 140502 (2021)
[2] S. Ahmed et al., Phys. Rev. Res. 3, 033278 (2021)
[3] https://github.com/quantshah/qst-cgan
[4] S. Ahmed et al., Phys. Rev. Lett. 130, 150402 (2023)
[5] https://github.com/quantshah/gd-qpt
[6] A. Gaikwad et al., in preparation (2025)

Neural matching decoder

Mats Granath

University of Gothenburg

We present a hybrid decoder for matchable quantum error correcting codes based on minimum weight perfect matching on a graph with machine learning generated edge weights. With large-scale experiments on quantum error correcting codes, such as the surface code, the availability of experimental data motivates the use of data-driven approaches to decoding. The approach combines an efficient and accurate graph algorithm-based decoder with an experimentally informed machine learning-based graph generation algorithm. The hybrid decoder is an alternative to standard approaches that involve separately deducing good edge weights by analysing the correlations between measured syndrome detectors, effectively combining this analysis with the decoding into a single algorithm. We benchmark the decoder against matching, belief-propagation augmented matching, and a pure machine learning-based decoder using simulated data.

Learning the scaling of multipartite entanglement quantification

Miroslav Jezek

Palacky University Olomouc

Entanglement is a fundamental resource for quantum technologies. Yet, the measurement cost of quantifying multipartite entanglement in large systems remains poorly understood. We explore the expressive power of deep neural networks to estimate the number of measurements required for entanglement quantification. Specifically, we use local single-copy measurements as inputs to a deep learning model trained to predict a target entanglement characteristic. For a fixed acceptable error in entanglement estimation, we identify the corresponding model and infer the minimal number of measurements required. By repeating this process for systems with an increasing number of particles (qubits), we aim to determine the scaling behavior of the measurement cost. Additionally, we discuss the potential for direct photonic entanglement quantification using random optical networks.

Do We Really Need Symmetry Functions to Understand Phase Transitions in Classical Systems?

Carina Kamer

TU Wien

Phase transitions in classical systems—such as the crystallization of a supercooled liquid—are inherently complex as they rely on the collective rearrangement from one phase into another. Although the full information of a system of N particles is, in principle, contained in a 6N-dimensional phase space of positions and momenta, reasoning in this high-dimensional space remains inaccessible due to our limited spatial imagination. As a result, statistical mechanics has long relied on phenomenological theories and dimensionality reduction via collective variables to make sense of phase transition dynamics.

One such approach is classical nucleation theory (CNT), which tracks crystallization by monitoring the growth of the largest crystalline cluster. This method depends critically on symmetry functions that detect local crystallinity, allowing for the identification of solid-like particles. While CNT performs reasonably well for simple model systems such as Lennard-Jones fluids, it completely fails to capture the full complexity of phase transitions in more realistic or biologically relevant systems, such as water or biominerals. But even for the simple Lennard-Jones system, open questions, such as the existence of a multi-step nucleation remain [1].

With the rise of machine learning, new opportunities have emerged to revisit these high-dimensional spaces. In particular, autoencoders—neural networks that learn compressed latent representations of input data—offer a new route to uncover structure in configuration space without pre-imposing physical assumptions. So far, however, autoencoder-based studies of crystallization still rely on symmetry functions, justified by the idea that phase transitions are ultimately about the emergence of crystalline symmetry rather than individual particle positions [2].

In this work, we challenge this assumption by exploring whether latent representations of particle trajectories during crystallization retain meaningful structure without any symmetry-based input. Can autoencoders uncover phase transition signatures directly from raw configurations, or is the signal lost in noise? This study addresses this long-standing question in computational physics and may offer new insights into how much of our current understanding depends on physically motivated but potentially limiting prior assumptions.

[1] Analyzing multistep homogeneous nucleation in vapor-to-solid transitions using molecular dynamics simulations, K.K. Tanaka, J. Diemand, H. Tanaka, R. Angelil, Phys. Rev. E 96, 022804 (2017)

[2] Autonomously revealing hidden local structures in supercooled liquid, E. Boattini, S. Marjin-Aguilar, S. Mitra, G. Foffi, F. Smallenburg and L. Filion, Nature Communications, volume 11, Article number: 5479 (2020)

Cross-Platform Autonomous Control of Minimal Kitaev Chains with Neural Networks

Rouven Koch

TU Delft / QuTech

A one-dimensional Kitaev model implemented in a quantum dot system coupled to superconductors can support Majorana zero modes at the ends of the chain. Recent experiments [1] have shown that even a two-dot system can host “Poor Man’s Majorana” modes in a specific regime of the Hamiltonian, known as the sweet spot [2]. In this regime, the strengths of elastic co-tunneling (ECT) and crossed Andreev reflection (CAR) are equal, which is crucial for the appearance of Poor Man’s Majoranas. Previous studies have demonstrated that a generative machine learning model can predict the underlying Hamiltonian parameters based on experimental data [3]. In this work, we introduce an automated tuning algorithm that leverages a convolutional neural network to infer the Hamiltonian state, allowing the quantum dot system to reach the sweet spot regime autonomously [4]. By combining theoretical insights, machine learning, and experimental techniques, we lay the groundwork for efficient, automated tuning of longer Kitaev chains, with promising applications in quantum information and computing.

[1] Tom Dvir, et al. "Realization of a minimal Kitaev chain in coupled quantum dots." Nature 614.7948 (2023): 445-450.
[2] Martin Leijnse, and Karsten Flensberg. "Parity qubits and poor man's Majorana bound states in double quantum dots." Physical Review B 86.13 (2012): 134528.
[3] Rouven Koch, David van Driel, Alberto Bordin, Jose L. Lado, and Eliska Greplova. “Adversarial Hamiltonian learning of quantum dots in a minimal Kitaev chain.” Phys. Rev. Applied 20, 044081 (2023).
[4] David van Driel, et al. "Cross-Platform Autonomous Control of Minimal Kitaev Chains." arXiv preprint arXiv:2405.04596 (2024).

Distributed quantum dynamics on near-term quantum processors

Maciej Koch-Janusz

Haiqu, Inc. / University of Zurich

TBA.

TBA

Roman Krems

University of British Columbia

TBA.

A relativistic variational quantum circuit

Marius Krumm

University of Innsbruck

The field of relativistic quantum information seeks to understand the quantum information properties of relativistic quantum systems. A popular approach for this purpose is the Unruh-DeWitt model for qubits interacting with quantum fields on curved spacetime. In my talk, I will present a relativistic variational quantum circuit (VQC) in which the interaction between qubits is mediated by the relativistic quantum field. An important consequence is that the tunable time evolution of the qubits depends on spacetime properties and quantum field propagators. Therefore, our VQC presents first steps in a quantum machine learning approach that seeks to extract quantum properties of spacetime and fields when no hand-crafted protocol is available. Our approach works in a regime in which the time evolution of the qubits is unitary and the quantum field does not act as a decohering environment, which we believe to be interesting in its own right.

Efficient quantum-enhanced classical simulation for patches of quantum landscapes

Sacha Lerch

EPFL

Understanding the capabilities of classical simulation methods is key to identifying where quantum computers are advantageous. Not only does this ensure that quantum computers are used only where necessary, but also one can potentially identify subroutines that can be offloaded onto a classical device. In this work, we show that it is always possible to generate a classical surrogate of a sub-region (dubbed a "patch") of an expectation landscape produced by a parameterized quantum circuit. That is, we provide a quantum-enhanced classical algorithm which, after simple measurements on a quantum device, allows one to classically simulate approximate expectation values of a subregion of a landscape. We provide time and sample complexity guarantees for a range of families of circuits of interest, and further numerically demonstrate our simulation algorithms on an exactly verifiable simulation of a Hamiltonian variational ansatz and long-time dynamics simulation on a 127-qubit heavy-hex topology.

Hamiltonian learning quantum magnets with non-local impurity tomography

Greta Lupi

EPFL

Impurities in quantum materials have provided successful strategies for learning properties of complex states, ranging from unconventional superconductors to topological insulators. In quantum magnetism, inferring the Hamiltonian of an engineered system becomes a challenging open problem in the presence of complex interactions. Here we show how a supervised machine-learning technique can be used to infer Hamiltonian parameters from atomically engineered quantum magnets by inferring fluctuations of the ground states due to the presence of impurities. We demonstrate our methodology both with a fermionic model with spin-orbit coupling, as well as with many-body spin models with long-range exchange and anisotropic exchange interactions. We show that our approach enables performing Hamiltonian extraction in the presence of significant noise, providing a strategy to perform Hamiltonian learning with experimental observables in atomic-scale quantum magnets. Our results establish a strategy to perform Hamiltonian learning by exploiting the impact of impurities in complex quantum many-body states.

Discussion session: Large-language models – what could they be good for in science?

Florian Marquardt

Max Planck Institute for the Science of Light

I will try to guide a discussion/brainstorming session on the topic of LLMs for science.

Out of the loop: Eliminating parameter optimisation in quantum-classical heuristics

Wolfgang Mauerer

Technical University of Applied Science Regensburg

Many heuristic quantum algorithms employ a quantum-classical iterative structure that intersperses the evaluation of quantum circuits with classical parameter optimisation phases. Unfortunately, these classical procedures are often known to be NP-complete, which somewhat counters the hope for finding quantum improvements to problems in this (and related) complexity classes.

Still, advantage over classical approaches is suspected for certain scenarios, but nature and origin of its computational power are not yet satisfactorily understood. We discuss how a mix of computer science and physics can help to better understand the actual computation in hybrid heuristics, and how the need to compute classical parameters in each iteration can be avoided by pre-determining a "good" set of parameters for a range of problem instances upfront.

TBA

Alexey Melnikov

Terra Quantum

TBA.

Quantum Reservoir Computing based on memristors

Daniel Montesinos Capacete

Institute for Cross-Disciplinary Physics and Complex Systems (IFISC)

We study a photonic reservoir computing framework inspired by the photonic quantum memristor design by Michele Spagnolo et al. (2022). Our design significantly enhances memory retention and non-linear dynamic capabilities by integrating an additional optical mode and using dual-rail encoding to feed back previous outputs into the system. Furthermore, starting with single memristors, we scale the framework via spatial multiplexing, employing random masks and diverse dynamic functions for each unit. This architecture notably improves short-term memory, enabling higher accuracy in complex nonlinear tasks such as autonomous forecasting of the Lorenz system and MNIST classification, showcasing its potential in quantum reservoir computing.

Representation learning reaches the lab: let machines act!

Gorka Muñoz-Gil

Innsbruck University

Scientific discovery is often linked to the extraction of a few key parameters that efficiently and accurately represent a physical process. Among the myriad machine learning approaches available, variational autoencoders have perhaps captured the most attention. These methods generally assume the existence of a proper dataset from which to extract relevant parameters. However, scientific discovery can occur even earlier: during the data collection phase in the laboratory! In this contribution, I will demonstrate how allowing a machine to interact freely with an experimental setup—through predefined actions, such as activating a magnetic field, measuring certain observables of a quantum state or even introducing a mutant in a cell—can enhance the extraction of relevant parameters. Moreover, this approach can also help us elucidate the relationships between these actions and the experiment's degrees of freedom.

Exploring topological protection in finite systems with interpretable machine learning

Dmytro Oriekhov

QuTech, TU Delft

A number of recent experiments with engineered quantum devices, for instance quantum dots and superconducting circuits, demonstrated a possibility to make one-dimensional topological insulators, like the SSH model. At the same time, the size of systems that has been demonstrated so far is limited due to engineering challenges connected to their experimental realization. That poses a question whether it is possible to obtain topological protection in these finite-size systems. In this contribution, I explain two new approaches to studying topology in real-space, finite size, disordered systems: (1) Detailed analysis of the finite size effect breaking topological protection due to the cross-talk between edge states. Leveraging this understanding, we deploy bulk conductivity as a candidate for experimental measurement that helps evaluate the remaining topological projection and analyze its behavior in the presence of disorder. In (2) we present a hybrid combination of neural and tensor networks that accurately represent the topology of the system and thus, is not susceptible to common misclassifications that result from noise or disorder. Combining approaches (1) and (2) creates a new framework for characterization of the contemporary and near-term finite-size topological experiments.

Guided-SPSA: Simultaneous Perturbation Stochastic Approximation assisted by the Parameter Shift Rule

Maniraman Periyasamy

Frauhofer-IIS / OTH Regensburg

The computational complexity, in terms of the number of circuit evaluations required for gradient estimation by the parameter-shift rule, scales linearly with the number of parameters in VQCs. On the other hand, techniques that approximate the gradients of the VQCs, such as the simultaneous perturbation stochastic approximation (SPSA), do not scale with the number of parameters but struggle with instability and often attain suboptimal solutions. In this talk, we introduce a gradient estimation approach called Guided-SPSA, which meaningfully combines the parameter-shift rule and SPSA-based gradient approximation. The Guided-SPSA results in a 15% to 25% reduction in the number of circuit evaluations required during training for a similar or better optimality of the solution found compared to the parameter-shift rule.

Guarantees and limitations for warm starts and iterative methods in variational quantum computing

Ricard Puig

EPFL

Barren plateaus are fundamentally a statement about quantum loss landscapes on average but there can exist patches of barren plateau landscapes with substantial gradients. This has motivated the study of warm starts whereby the algorithm is cleverly initialized closer to a minimum. Numerical studies indicate that these methods may be promising. In parallel, analytic studies have proven that small angle initializations, whereby the parameterized quantum circuit is initialized in a small region typically around identity or a Clifford, can exhibit non-exponentially vanishing gradients. However, a good solution may be far from this region and thus these methods can (in full generality) only work on a vanishing fraction of problem instances. In this joint submission we present general analysis of warm starts for physically-motivated ansatze and iterative training strategies. Our work thus suggests that while there are hopes to be able to warm-start variational quantum algorithms, any initialization strategy that cannot get increasingly close to the region of attraction with increasing problem size is likely to prove challenging to train.

https://arxiv.org/abs/2502.07889

https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.6.010317

QUANDELA

Neural networks leverage nominally quantum and post-quantum representations

Paul Riechers

Simplex; Beyond Institute for Theoretical Science

We show that deep neural networks, like transformers and RNNs, pretrained as usual on next-token prediction, intrinsically discover and represent beliefs over 'quantum' and 'post-quantum' low-dimensional models of classical stochastic processes. We anticipate and find specific architecture-independent geometric relationships among activations induced by different inputs. The geometric structure corresponds to the correlational structure of whatever stochastic process generates the data. The points in this geometry correspond to history-induced probability densities over all possible futures.

Neural Quantum States as Dynamical Mean Field Theory Solvers

Jonas Rigo

Forschungszentrum Jülich

Neural Quantum Sates (NQS) constitute a variational wave function ansatz, that can provably efficiently represent even highly entangled quantum many-body states. Beyond their representative power, NQS inherit the speed of modern neural networks (NN) and equally profit from the enormous development that NNs have recently received. In this work we show that NQS can efficiently find the ground state of quantum impurity models with large baths, allowing us to compute high quality real-frequency, zero-temperature Green's functions by means of a Krylov-like method. We demonstrate the capability of this approach and its potential as dynamical mean-field theory (DMFT) solver at the example of the Bethe lattice and other benchmarks.

Improving Quantum Machine Learning via Heat-Bath Algorithmic Cooling

Nayeli Azucena Rodriguez Briones

TU Wien Atominstitut

In this talk, I will present a novel approach rooted in quantum thermodynamics to improve sampling efficiency in quantum machine learning (QML). The key idea is to conceptualize quantum supervised learning as a thermodynamic cooling process. Based on this perspective, we introduce a quantum refrigerator protocol that enhances sample efficiency during both training and prediction, without relying on Grover iterations or quantum phase estimation. Inspired by heat-bath algorithmic cooling protocols, our method leverages alternating entropy compression and thermalization steps to reduce qubit entropy, thereby increasing polarization towards the dominant bias. This technique minimizes the computational overhead associated with estimating classification scores and gradients, making it a practical and efficient solution for QML algorithms on noisy intermediate-scale quantum (NISQ) devices.

Reference: NA Rodríguez-Briones, DK Park. arXiv preprint arXiv:2501.02687

TBA

Asel Sagingalieva

Terra Quantum

TBA.

Application of ML to Many-Body quantum experiments; An experimental view

Jörg Schmiedmayer

TU Wien Atominstitut

In my talk I will give an overview how we as an experimental group use machine learning to control, read and analyse our many body quantum experiments.

Regarding experimental control we applied Physics-Inspired Learning Algorithms to optimize Optical Potentials and potential corrections for ultra cold quantum gases and quantum fields [1]. Tis gave us a speedup of nearly a factor 10 in adapting and verifying the light-fields so that the atoms see the desired potential landscape.

Regarding readout and measurement, we employed neural networks to significantly improve our methods to estimate the temperature of our quantum degenerate atomic ensembles [2]. The neural networks were able to estimate the temperature of the quantum degenerate atomic ensemble from the fluctuations, i.e. the matter wave spackle pattern, from much fewer experiments.

Finally, I will discuss how to extract transport properties in quantum gases. In order to obtain both the atomic and energy currents from the sparsely measured atomic density, we employ a Physics Informed Neural Network (PINN) [3]. By ensuring consistency with the conservation laws, an enhanced fidelity of the reconstructed currents compared to methods based on finite difference can be achieved.

Work performed in collaboration with the groups of A. Kugi and A. Deutschman (TU-Wien) and T. Calarco (Jülich). Supported by the DFG-FWF SFB ISOQUANT, and the ERC-AdG *Emergence in Quantum Physcs (EmQ)*

[1] Optimizing optical potentials with physics-inspired learning algorithms
M. Calzavara, et al. Phys. Rev. Applied **19**, 044090 (2023) arXiv:2210.07776

[2] Thermometry of one-dimensional Bose gases with neural networks
F. Møller, et al. Phys. Rev. A **104**, 043305 (2021), arXiv:2105.03127

[3] Characterising transport in a quantum gas by measuring Drude weights
P. Schüttelkopf, et al. arXiv:2406.17569

Quantum resources of quantum and classical variational methods

Thomas Spriggs

TU Delft

Variational techniques have long been at the heart of atomic, solid-state, and many-body physics. They have recently extended to quantum and classical machine learning, providing a basis for representing quantum states via neural networks. These methods generally aim to minimize the energy of a given ansätz, though open questions remain about the expressivity of quantum and classical variational ansätze. The connection between variational techniques and quantum computing, through variational quantum algorithms, offers opportunities to explore the quantum complexity of classical methods. We demonstrate how the concept of non-stabilizerness, or magic, can create a bridge between quantum information and variational techniques and we show that energy accuracy is a necessary but not always sufficient condition for accuracy in non-stabilizerness. Through systematic benchmarking of neural network quantum states, matrix product states, and variational quantum methods, we show that while classical techniques are more accurate in non-stabilizerness, not accounting for the symmetries of the system can have a severe impact on this accuracy. Our findings form a basis for a universal expressivity characterization of both quantum and classical variational methods.

Of Molecules and Stars: Device-agnostic Approach to Super-Resolution Imaging

Dominik Vašinka

Palacký University Olomouc

Measurement correlations in quantum systems can exhibit non-local behavior, a fundamental aspect of quantum mechanics with applications such as device-independent quantum information processing. However, it is in general not known which states are local and which ones are not. In particular, it remains an outstanding challenge to explicitly construct local hidden-variable (LHV) models for arbitrary multipartite entangled states. To address this, we use gradient-descent algorithms from machine learning to find LHV models which reproduce the statistics of arbitrary measurements for quantum many-body states. In contrast to previous approaches, our method employs a general ansatz, enabling it to discover LHV models for all local states. Therefore, it for example provides actual estimates for the critical noise levels at which two-qubit Werner states and three-qubit GHZ and W states become local. Furthermore, we find evidence suggesting that two-spin subsystems in the ground states of translationally invariant Hamiltonians are genuinely local, while bigger subsystems are in general not. Our method now offers a quantitative tool for determining the regimes of non-locality in any given physical context, such as non-equilibrium, decoherence or disorder.

Discovering Local Hidden-Variable Models for Arbitrary Multipartite Entangled States and Arbitrary Measurements

Nick von Selzam

Max Planck Institute for the Science of Light

Measurement correlations in quantum systems can exhibit non-local behavior, a fundamental aspect of quantum mechanics with applications such as device-independent quantum information processing. However, it is in general not known which states are local and which ones are not. In particular, it remains an outstanding challenge to explicitly construct local hidden-variable (LHV) models for arbitrary multipartite entangled states. To address this, we use gradient-descent algorithms from machine learning to find LHV models which reproduce the statistics of arbitrary measurements for quantum many-body states. In contrast to previous approaches, our method employs a general ansatz, enabling it to discover LHV models for all local states. Therefore, it for example provides actual estimates for the critical noise levels at which two-qubit Werner states and three-qubit GHZ and W states become local. Furthermore, we find evidence suggesting that two-spin subsystems in the ground states of translationally invariant Hamiltonians are genuinely local, while bigger subsystems are in general not. Our method now offers a quantitative tool for determining the regimes of non-locality in any given physical context, such as non-equilibrium, decoherence or disorder.

Problem-informed Graphical Quantum Generative Learning

Zoltan Zimboras

Wigner Research Centre for Physics, Hungary

Leveraging the intrinsic probabilistic nature of quantum systems, generative quantum machine learning (QML) offers the potential to outperform classical learning models. Current generative QML algorithms mostly rely on general-purpose models that, while being very expressive, face several training challenges. A potential way to address these setbacks involves constructing problem-informed models capable of more efficient training on structured problems. In particular, probabilistic graphical models provide a flexible framework for representing structure in generative learning problems and can thus be exploited to incorporate inductive bias in QML algorithms. In this work, we propose a problem-informed quantum circuit Born machine Ansatz for learning the joint probability distribution of random variables, with independence relations efficiently represented by a Markov network (MN). We further demonstrate the applicability of the MN framework in constructing generative learning benchmarks and compare our model's performance to previous designs, showing it outperforms problem-agnostic circuits. Based on a preliminary analysis of trainability, we narrow down the class of MNs to those exhibiting favorable trainability properties. Finally, we discuss the potential of our model to offer quantum advantage in the context of generative learning.

Hardness of sampling Gibbs states

Alexander Zlokapa

MIT

Classically, Metropolis-like algorithms such as MCMC are central to sampling tasks in learning and beyond. A recent line of work has introduced analogous quantum algorithms, but their associated mixing time is poorly understood. We provide sufficient conditions for the slow mixing of these quantum algorithms, allowing us to show (1) classical Hamiltonians that are hard for classical MCMC remain hard for these quantum algorithms, and (2) hardness of sampling from Gibbs states of non-commuting Hamiltonians, such as the 2D antiferromagnetic transverse field Ising model.

Posters

Entanglement transitions in quantum games

Giovanni Cemin

Max Planck Institute for the Physics of Complex Systems

In this work, we explore entanglement dynamics in gamified Clifford circuits by employing a reinforcement learning (RL) agent competing against a random agent. The RL agent strategically places gates to reduce entanglement, while the random agent increases it. This competitive setup induces an entanglement transition, whose characteristics depend on the amount of information available to the RL agent. By systematically varying the system size, the relative frequency of moves executed by the two agents, and the information provided to the RL agent, we investigate their effects on the entanglement transition. Our results reveal new insights into the interplay between entanglement control and information constraints, offering a deeper understanding of the mechanisms driving quantum circuit dynamics.

Dynamically disentangled state-action representation learning

Arunava Majumder

University of Innsbruck

To understand the behavior of a physical system or phenomenon, we often seek a minimal abstract representation that accurately describes it. Such representations generally focus on the most critical features or underlying factors of variation that are sufficient for their description while discarding redundant details.

In disentangled representation learning, models are designed to identify and separate these underlying factors hidden within the observed data. When the data originate from a physical system, the hidden factors typically correspond to the system's degrees of freedom (e.g., the mass or charge of an object). Recently, the ability of Variational Autoencoders (VAEs) to learn such representations has rendered them particularly appealing for practical applications.

In order to achieve a minimal yet informative representation of a physical system, scientific experiments consist of actions or observations, such as measuring observables on a quantum system or activating a magnetic field that captures the system's true characteristics. In this work, we introduce a variant of the VAE that extracts the relevant parameters by incorporating both the resultant observed data and the associated actions performed on the system. By design, this model architecture elucidates the relationship between these actions and the corresponding degrees of freedom, thereby enhancing the model's overall explainability for general-purpose tasks.