MAX PLANCK INSTITUTE
FOR THE SCIENCE OF LIGHT

Munich Quantum Valley

2nd Workshop of Machine Learning for Quantum Technology

November 5th—November 8th, 2024

Max Planck Institute for the Science of Light, Erlangen, Germany

Book of abstracts

Organizers

Florian Marquardt
Jan Olle
Matteo Puviani
Oleg Yevtushenko
Remmy Zen
Sangkha Borah
Gesine Murphy

Condensed Matter Physics with Machine Learning

Eliška Greplová
Delft University of Technology

In this tutorial, we are going to walk through some first historic (ca 2016!) applications of ML in condensed matter physics. We will look at how to learn phase transitions in supervised and unsupervised way and how the inheritance of these first ML+quantum papers propagates to how the field looks today.

Reinforcement Learning for Quantum Technologies

Marín Bukov
MPI-PKS

abstract

Autoregressive Models in Quantum Physics

Roger Melko
University of Waterloo

abstract

Autonomous Quantum Control in the Age of AI

Eliška Greplová
Delft University of Technology

Contemporary quantum computers and simulators keep increasing in size and complexity. While this enormous progress allows both academia and industry to explore quantum systems at an unprecedented scale, the newest generation of devices also comes with unparalleled control and calibration challenges. Enter artificial intelligence: At the heart of all control problems lies a data interpretation task. How much can the predictive and generalization power of AI help us tackle these challenges? In this talk I will explain how to apply modern computational methods, including AI, to control and learn physics of contemporary quantum devices and I will give a perspective on how AI can aid us as we scale up the newest generation of quantum technologies.

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 first 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]. Finally, I will present recent results applying some of these ideas that we used for process tomography to state tomography [4]. Our codes for both state and process tomography are freely available on GitHub [5].

[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] A. Gaikwad et al., in preparation (2024)
[5] https://github.com/quantshah/qst-cgan; https://github.com/quantshah/gd-qpt

Machine learning assisted quantum simulator readout

Martin Gärttner
Friedrich-Schiller-University Jena

Quantum simulation seeks to emulate hard to solve quantum many-body problems using synthetic quantum systems. One challenge in this endeavour is the resource efficient characterization of the prepared states and the extraction of observables of interest. We leverage machine learning methods, including neural quantum state, supervised learning methods, and spiking neuromorphic hardware to alleviate this challenge.

Realizing a reinforcement learning agent for real-time quantum feedback

Christopher Eichler
FAU Erlangen

Rapid advancements in building quantum information processing devices at scale call for radical paradigm shifts in the real-time control of quantum systems. Reinforcement learning offers a powerful approach to learn sophisticated control strategies in the absence of detailed models and promises to become a game changer in quantum technology as it did in many other disciplines ranging from board games, to robotics, to fundamental science. In my talk, I will show how we achieve real-time feedback control of a quantum system by using a reinforcement learning agent. We realize the agent as a novel low-latency neural-network on a field-programmable gate array interacting with a superconducting quantum system at MHz rates, which is more than 100 times faster than any other previous implementation of a reinforcement learning agent deployed in any physics experiment. Our work paves the way towards using reinforcement learning for real-time control of quantum computers, most notably for quantum error correction and fault-tolerant gate operations. Beyond that, the unprecedented speed of the agent marks an important engineering achievement , which can be deployed in a wide range of other applications, such as real-time the control of optical systems.

K. Reuer, J. Landgraf, T. Fösel, J. O'Sullivan, L. Beltrán, A. Akin, G.J. Norris, A. Remm, J.-C. Besse, F. Marquardt, A. Wallraff and C. Eichler, Nat. Comm., 14:7138 (2023)

Trying to solve quantum many-body problems with neural networks

Annabelle Bohrdt
Universität Regensburg

Strongly correlated quantum many-body systems are often both hard to understand theoretically and hard to simulate numerically. In this talk, I will show how we can use neural networks to advance the current understanding of interacting quantum many-body systems. I will in particular focus on models relevant to strongly correlated electrons, starting by briefly introducing the relevant physics, and show how we can use neural networks to (i) represent ground states of t-J models, highly relevant in the context of unconventional superconductors; and (ii) analyze data, e.g. from quantum simulators, to determine relevant observables.

Physics-inspired machine learning models and optimal control for quantum experiments

Maximilian Prüfer
Atominstitut, TU Wien

Control of Hamiltonian parameters and initial conditions is pivotal for performing quantum simulations. We developed a physics-inspired model approach to optimize optical potentials for quantum experiments. The trained model, together with a feedback routine, allows us to perform offline optimization, that is, without using the experiment, with accuracies comparable to existing online methods. Further, we explore optimal control strategies to produce highly squeezed states in a tunnel-coupled double well with ultracold atoms. Our approach will lead to states with lower effective temperatures in our system, which simulates the sine-Gordon quantum field theory.

Error-tolerant quantum convolutional neural networks for symmetry-protected topological phases

Petr Zapletal
University of Basel

The analysis of noisy quantum states prepared on current quantum computers is getting beyond the capabilities of classical computing. Quantum neural networks based on parametrized quantum circuits, measurements and feed-forward can process large amounts of quantum data to reduce measurement and computational costs of detecting nonlocal quantum correlations. Characterizing non-local correlations is crucial in condensed matter physics for classifying quantum phases of matter and understanding strongly correlated materials such as topological quantum matter. A key requirement for employing quantum neural networks to characterize noisy quantum data produced by near-term quantum hardware is tolerance to errors due to decoherence and gate infidelities. In Ref. [1], we construct quantum convolutional neural networks (QCNNs) that can, in the presence of incoherent errors, recognize different symmetry-protected topological phases of generalized cluster-Ising Hamiltonians from one another as well as from topologically trivial phases. Using matrix product state simulations, we show that the QCNN output is robust against symmetry-breaking errors below a threshold error probability and against symmetry-preserving errors provided the error channel is invertible. This is in contrast to string order parameters and the output of previously designed QCNNs, which vanish in the presence of any symmetry-breaking errors. The QCNN circuits can be shortened from logarithmic to constant depth in system size by performing a large part of the computation in classical postprocessing. This allowed us to realize the error-tolerant QCNNs on a 7-qubit superconducting quantum processor [2]. These constant-depth QCNNs reduce sample complexity exponentially with system size in comparison to the direct sampling using local Pauli measurements.

[1] P. Zapletal, N. A. McMahon, M. J. Hartmann, Phys. Rev. Research 6, 033111 (2024)
[2] J. Herrmann et al., Nat. Commun. 13, 4144 (2022)

Reinforcement learning transmon-qubit entangling gates

Marín Bukov
MPI-PKS

The utility of a quantum computer is highly dependent on the ability to reliably perform accurate quantum logic operations. For finding optimal control solutions, it is of particular interest to explore model-free approaches, since their quality is not constrained by the limited accuracy of theoretical models for the quantum processor -- in contrast to many established gate implementation strategies.
In this work, we utilize a continuous control reinforcement learning algorithm to design entangling two-qubit gates for superconducting qubits; specifically, our agent constructs cross-resonance and CNOT gates without any prior information about the physical system.
Using a simulated environment of fixed-frequency fixed-coupling transmon qubits, we demonstrate the capability to generate novel pulse sequences that outperform the standard cross-resonance gates in both fidelity and gate duration, while maintaining a comparable susceptibility to stochastic unitary noise. We further showcase an augmentation in training and input information that allows our agent to adapt its pulse design abilities to drifting hardware characteristics, importantly, with little to no additional optimization.
Our results exhibit clearly the advantages of unbiased adaptive-feedback learning-based optimization methods for transmon gate design.

Mach. Learn.: Sci. Technol. 5 (2024) 025066

Calibration of decoders for quantum error correction using multi-agent reinforcement learning

Volodymyr Sivak
Google Quantum AI

Accurate decoding of quantum error-correcting codes is a crucial ingredient in protecting quantum information from decoherence. It requires characterizing the error channels corrupting the logical quantum state and providing this information as a prior to the decoder, a process known as decoder calibration. In Ref. [1], we introduced a calibration method inspired by multi-agent reinforcement learning, that aims to minimize the logical error rate. It was recently used for calibrating decoders in an experiment that demonstrated below-threshold error correction with the surface code in Ref. [2]. This calibration approach will serve as an important tool for maximizing the performance of both near-term and future error-corrected quantum devices.

[1] V. Sivak, M. Newman, P. Klimov, arXiv:2406.02700 (2024)
[2] Google Quantum AI, arXiv:2408.13687 (2024)

Analyzing and constructing efficient data encoding quantum circuits

Maciej Koch-Janusz
Haiqu, Inc.

Encoding classical data in a quantum state is a key prerequisite of many quantum algorithms. Recently tensor methods emerged as the most promising approach for constructing efficient quantum circuits approximating input distributions. We derive rigorous bounds for the decay of entanglement across bonds in a tensorial representation of a function, real or complex, depending on its smoothness and localization properties. Based on these analytical results we construct an improved algorithm resulting in shallow and accurate encoding quantum circuits. We validate our methods on heavy-tailed distributions, important in finance, and in particular on Levy distributions.

Fast Hardware-efficient Quantum Gate Design using Optimal Control with Reinforcement Learning Ansatz

Bijita Sarma
FAU Erlangen-Nuremberg

Designing fast quantum gates is crucial for the prospect of achieving fault-tolerant quantum computation as well as for current noisy quantum hardware so that the detrimental effects of decoherence can be minimised during the operation of the gates. However, achieving fast gates with high-fidelity and desirable efficiency on the state-of-the-art physical hardware platforms remains a formidable task owing to the presence of hardware-level leakage errors and crosstalk. In recent years, machine learning (ML)-based methods have found widespread applications in different domains of science and technology for nontrivial tasks. In this work, we exploit the power of ML, in particular of reinforcement learning (RL), to design fast two-qubit quantum gates, viz. a controlled-Z gate on multi-level transmon qubits in a tunable coupler circuit. We show that the RL-controlled pulses can reduce hardware-level leakage errors despite the gate being fast, resulting in high fidelity. Additionally, we combine gradient-based optimal control (OC) with RL-optimization, that results in a superior optimizer compared to using RL or OC independently. This improved optimizer achieves a fidelity above 99.99% within a 20 ns gate time.

Graph neural network based decoders for quantum error correcting codes

Mats Granath
University of Gotehnburg

Quantum error correcting codes protect logical information by encoding it in an extended state over many data-qubits. Scaling up the size of the code, with overall error rate below a certain threshold, gives an exponential enhancement of the logical lifetime. Crucial for the performance is a classical decoder algorithm that interprets sequences of stabilizer measurements to suggest the best error correction operation. Recently, this scaling has been shown experimentally using surface codes with 9, 25, and 49 data-qubits [1]. In these recent experiments the best decoder, giving the lowest logical failure rates, is a neural network (NN) decoder. Surprisingly, this decoder even outperforms a tensor-network-based maximum-likelihood decoder (MLD). The advantage of the NN decoder can be traced to the fact that it is model-free, trained on experimental data, whereas the MLD requires an over-simplified error model.

I will present our work on a relatively simple but powerful decoder using graph neural networks (GNNs) [2]. 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 belief-propagation augmented matching decoder for circuit level noise on the surface code for simulated experimental data. Work in progress that may be discussed include decoding low-density parity check codes, using the GNN decoder as a soft-output decoder, implementing the decoder on FPGAs for fast inference, and efforts to build a hybrid “neural matching” decoder.

[1] Google Quantum AI and Collaborators, arXiv2408.13687
[2] Lange et al., arXiv:2307.01241

Quantum many-body systems under the microscope

Monika Aidelsburger
MPQ/LMU

abstract

Decoding many-body teleportation

Simon Trebst
University of Cologne

Teleportation is a facet where quantum measurements can act as a powerful resource in quantum physics, as local measurements allow to steer quantum information in a non-local way. While this has long been established for a single Bell pair, the teleportation of a many-qubit entangled state using non-maximally entangled resources presents a fundamentally different challenge. In this talk I will discuss a tangible protocol for teleporting a long-range entangled surface code state using elementary Bell measurements and its stability in the presence of coherent errors that weaken the Bell entanglement. Notably, this teleportation protocol points to new venues for quantum error correction allowing to push the threshold for successful teleportation.

Neural-network-assisted parameter estimation for quantum detection

Yue Ban
Universidad Carlos III de Madrid

Quantum sensing and quantum metrology have experienced significant progress, placing themselves at the forefront of the new generation of technologies. In this presentation, parameter estimation in different quantum platforms assisted by neural networks are introduced. Our results show that neural networks are valuable for quantum sensing leading to adaptive protocols for quantum detection with broad working regime and high accuracy. The benefits to integrate neural networks are illustrated to decipher the information contained in the sensor responses at the data processing stage of general quantum sensing tasks. We experimentally demonstrate that the combination of 171Yb+ atomic sensors with adequately trained neural networks enables to investigate target fields in scenarios where the responses go beyond their standard harmonic behavior [1, 2] in the presence of large shot noise. We also present a signal-to-image deep learning model capable of automatically inferring the number of nuclear spins surrounding a NV sensor and the hyperfine couplings between the sensor and the nuclear spins [3]. The way for the practical use of quantum many-body systems as black-box sensors using neural networks is also demonstrated. Neural networks faithfully reproduce the dynamics of quantum many-body sensors [4], thus allowing for an efficient Bayesian analysis.

References:
[1] Y. Chen, Y. Ban, R. He, J.-M. Cui, Y.-F. Huang, C.-F. Li, G.-C. Guo, J. Casanova, A neural network assisted 171Yb+ quantum magnetometer. Npj Quantum Inf. 8, 152 (2022).
[2] Y. Ban, J. Echanobe, Y. Ding, R. Puebla, J. Casanova, Neural-network-based parameter estimation for quantum detection, Quantum Sci. Technol. 6, 045012 (2021).
[3] B. Varona-Uriarte, C. Munuera-Javaloy, E. Terradillos, Y. Ban, A. Alvarez-Gila, E. Garrote, and J. Casanova, Automatic Detection of Nuclear Spins at Arbitrary Magnetic Fields via Signal-to-Image AI Model, Phys. Rev. Lett. 132, 150801 (2024).
[4] Y. Ban, J. Casanova, R. Puebla, Neural networks for Bayesian quantum many-body magnetometry, arXiv: 2212.12058 (2022).

Neural quantum states for many-body electronic structure and dynamics

Giuseppe Carleo
EPFL

This talk presents recent advances in using neural-network parameterizations of many-body wave functions for ab initio electronic structure calculations. We introduce a message-passing neural network ansatz for simulating strongly interacting electrons in continuous space, demonstrating its high accuracy for the homogeneous electron gas problem and for the unitary Fermi gas. The approach's potential for other periodic systems and materials will be discussed. Additionally, we explore the application of neural wave functions to the challenging case of many-body electronic quantum dynamics, showcasing results for quantum dots and molecules in strong laser fields.

(Neural) network representations of many-body wave functions

Markus Schmitt
Universität Regensburg

Modern quantum simulators are opening up the realm of non-equilibrium many-body physics in two spatial dimensions, where a rich phenomenology can be expected for example due to the existence of extended interfaces or phase transitions at non-vanishing energy densities. However, the numerical simulation of such situations, which provides the crucial link to theoretical models, remains a pivotal challenge. Moreover, it is an open question, how wave functions realized in a quantum simulator can be probed most comprehensively given restricted measurement budgets. In this context, I will discuss how neural quantum states (NQS) can be used to simulate the dynamics of Rydberg atom arrays. I will also introduce the framework of wave function networks as a new approach to characterize many-body states based on snapshot measurements beyond few-body correlations. These wave function networks can be constructed equivalently using data from NQS simulations or experiments, with the presence of correlations is reflected in their non-trivial structure. Furthermore, the network properties can serve as a basis for cross-certification between different sources.

Paths towards time evolution with larger neural-network quantum states

Dario Poletti
Singapore University of Technology and Design

In recent years, the neural-network quantum states method has been investigated to study the ground state and the time evolution of many-body quantum systems. Here we expand on the investigation and consider a quantum quench from the paramagnetic to the anti-ferromagnetic phase in the tilted Ising model. We use two types of neural networks, a restricted Boltzmann machine and a feed-forward neural network. We show that for both types of networks, the projected time-dependent variational Monte Carlo (p-tVMC) method performs better than the non-projected approach. We further demonstrate that one can use K-FAC or minSR in conjunction with p-tVMC to reduce the computational complexity of the stochastic reconfiguration approach, thus allowing the use of these techniques for neural networks with more parameters.

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.

Language Models for Quantum Simulation

Roger Melko
University of Waterloo

Generative language models are causing significant technological disruption, showing remarkable abilities on typical language tasks - such as generating text from a prompt. However, these model architectures are also extremely well-suited to many other generative tasks, and are being rapidly adopted as numerical simulation tools for quantum many-body systems. In this talk, I will discuss the use of language models for learning quantum states realized in today's experimental devices. I will argue that, if we can leverage their ability to scale, language models could become one of the most powerful computational tools in our arsenal for the design, characterization and control of tomorrow’s quantum computers.

Solving 2D quantum matter with neural quantum states

Markus Heyl
University of Augsburg

Neural quantum states have emerged as a novel promising numerical method to solve the quantum many-body problem both in and out of equilibrium.
In this talk I will highlight the recent progress in particular concerning correlated quantum matter in two spatial dimensions both for the equilibrium ground-state as well nonequilibrium real-time dynamics problem. For the calculations of ground states I will introduce the minimum-step stochastic reconfiguration method that reduces the optimization complexity by orders of magnitude. I will show that with this method we can now accurately train on unprecedentedly deep neural quantum states with millions of parameters allowing us to obtain the lowest variational energies as compared to existing numerical results for frustrated quantum magnets. Further, I will highlight the recent results on solving the real-time dynamics of correlated quantum matte rusing neural quantum states, which has allowed us to verify for instance for the first time the quantum Kibble-Zurek mechanism for interacting quantum many-body systems in two spatial dimensions.

Learning the dynamics of Markovian open quantum systems from experimental data

Cristian Bonato
Heriot-Watt University

Automating the extraction of physical models from experimental data in a reliable, robust, and interpretable form is an important challenge with far-reaching implications for both fundamental scientific discoveries and technological applications. Here we report an algorithm, based on Markov-chain Monte Carlo, to automatically construct Lindblad master equation models for Markovian open quantum systems, from experimental data. While the mater equation description for a given quantum system is not unique, our algorithm outputs a ranked list of model classes, each describing a possible physical explanation for the data. We benchmark the algorithm on experimental data on two optical quantum dots. In the ideal case, this sytem could either behave as independent emitters (when non resonant), or exhibit super-radiance. The latter physics results in a shortening of the emission lifetime and in an anti-bunching peak in the second-order correlation function. As the experimental data does not fit either of the two ideal limiting cases [Koong et al., Sci. Adv. 8, eabm8171 (2022)], we apply our algorithm to find potential explanations, exploring the role of dephasing and asymmetry. In conclusion, we believe our model learning algorithm is a useful tool to explore the physics of open quantum systems. In particular we envison the algorithm to suggest to the experimenter possible alternative models that they might not have thought about when designing the experiment, and to help them ensure that the proposed theoretical mechanisms are the only possible interpretation of the experimental data.

Program synthesis-driven quantum architecture search for optimal quantum circuit design in variational quantum algorithms

Akash Kundu
QTF Centre of Excellence, Department of Physics, University of Helsinki, Finland

Building quantum circuits that perform a given task is a notoriously difficult problem. Reinforcement learning has proven to be a powerful approach, but many limitations remain due to the exponential scaling of the space of possible operations with qubits. In this paper, we develop an algorithm that automatically learns composite gates and adds them as additional actions to the reinforcement learning agent to facilitate the search. We apply this idea to finding circuits that implement the ground state of a given quantum Hamiltonian, a well-known NP-hard challenge. In particular, we focus on the Transverse Field Ising Model (TFIM) as the understanding of the ground state of the TFIM is crucial for studying quantum phase transitions and critical behavior, serving as a benchmark for validating quantum algorithms and simulation techniques. We investigate whether we can find a better circuit for implementation on real hardware and whether this technique scales better to larger systems compared to a pure reinforcement learning approach. This research raises critical questions about the optimality of the gate decomposition, the scalability to larger qubit systems, and the complexity inherent in determining optimal rotation angles. The generality of the algorithm will allow future applications to other settings, including optimization tailored to specific real-world quantum platforms.

Towards an Artificial Muse for new ideas in Science

Mario Krenn
Max Planck Institute for the Science of Light

Artificial intelligence (AI) is a potentially disruptive tool for physics and science in general. One crucial question is how this technology can contribute at a conceptual level to help acquire new scientific understanding or inspire new surprising ideas. I will talk about how AI can be used as an artificial muse in physics, which suggests surprising and unconventional ideas and techniques that the human scientist can interpret, understand and generalize to its fullest potential.

[1] Krenn, Kottmann, Tischler, Aspuru-Guzik, Conceptual understanding through efficient automated design of quantum optical experiments. Physical Review X 11(3), 031044 (2021).

[2] Krenn, Pollice, Guo, Aldeghi, Cervera-Lierta, Friederich, Gomes, Häse, Jinich, Nigam, Yao, Aspuru-Guzik, On scientific understanding with artificial intelligence. Nature Reviews Physics 4, 761–769 (2022).

[3] Krenn et al., Forecasting the future of artificial intelligence with machine learning-based link prediction in an exponentially growing knowledge network, Nature Machine Intelligence 5, 1326 (2023)

[4] Gu, Krenn, Interesting Scientific Idea Generation Using Knowledge Graphs and LLMs: Evaluations with 100 Research Group Leaders. arXiv:2405.17044 (2024)

Provable exponential quantum advantages in learning from classical data

Vedran Dunjko
Leiden University

One of the key challenges of the quantum machine learning field is identifying learning problems where quantum learning algorithms can achieve a provable exponential advantage over classical learning algorithms. Recent works showed how to use certain less-discussed complexity theoretic assumptions that allow us to make formal proofs of quantum learning advantages stemming from common complexity-theoretic separations between classical computers, including for e.g. in quantum simulation. The fundamental ideas are using an intuitive fact: if the functions we are trying to learn are too complex for any poly-sized classical circuit to represent, then (trivially) classical learning is impossible. In this talk we will briefly reflect on these ideas and show two new results. First, we will show how this framework can be used to prove learning separations in physically meaningful problems, namely of learning observables.
Second, we will present ongoing work on a new stronger result (due to time constraints, we will be happy if we present the statement of the problem in an accessible way TBH), showing that classical/quantum learning separations are possible even if the learner is only required to identify (and not evaluate!) the function we are trying to learn.

This talk is based on : arXiv:2306.16028, arXiv:2405.02027 and a still unpublished work.

Towards explainable AI in quantum science

Hans Briegel
University of Innsbruck

I will review some of our research on the use and development of AI in quantum physics and beyond. This includes the model of Projective Simulation as an interpretable model for learning and decision-making with applications, e.g., in the design of quantum experiments, quantum error correction and robotic skill acquisition. I will review some of these applications and the interplay of concepts from quantum information and machine learning, more generally. The latter includes a recent proposal of variational measurement-based quantum computation for generative modelling and the use of diffusion models for quantum circuit design.

Unveiling quantum phase transitions from traps in variational quantum algorithms

Chenfeng Cao
Phasecraft

Understanding quantum phase transitions in physical systems is fundamental to characterize their behaviour at small temperatures. Achieving this requires both accessing good approximations to the ground state and identifying order parameters to distinguish different phases. Addressing these challenges, our work introduces a hybrid algorithm that combines quantum optimization with classical machine learning. This approach leverages the capability of near-term quantum computers to prepare locally trapped states through finite optimization. Specifically, we utilize LASSO for identifying conventional phase transitions and the Transformer model for topological transitions, applying these with a sliding window of Hamiltonian parameters to learn appropriate order parameters and estimate the critical points accurately. We verified the effectiveness of our method with numerical simulation and real-hardware experiments on Rigetti's Ankaa 9Q-1 quantum computer. Our protocol not only provides a robust framework for investigating quantum phase transitions using shallow quantum circuits but also significantly enhances efficiency and precision, opening new avenues in the integration of quantum computing and machine learning

Machine Learning for Fault-Tolerant Quantum Computation

Johannes Bausch
Google Deepmind

Fault-tolerant quantum computation requires both effective quantum error correction and minimization of resource overhead for practical applications. We present two examples of leveraging machine learning (ML) to address these challenges. First, we introduce a novel recurrent, transformer-based neural network that learns to decode the surface code. This ML decoder surpasses the accuracy of state-of-the-art algorithmic decoders on experimental data from Google's Sycamore processor and maintains this advantage at larger scales in simulations with realistic noise, including crosstalk and leakage. Second, we present AlphaTensor-Quantum, a deep reinforcement learning agent for optimizing quantum circuits by reducing the count of T gates, which remains a major contributor to the runtime and resource costs of fault-tolerant quantum algorithms. Exploiting the relationship between T-count optimization and tensor decomposition, and incorporating domain-specific knowledge, AlphaTensor-Quantum significantly outperforms existing methods on a benchmark of arithmetic circuits and discovers solutions matching human-designed implementations for relevant quantum computations, including those used in quantum chemistry simulations.

RL and RL for quantum systems

Evert van Nieuwenburg
Leiden University

In this talk I will focus on two interesting quantum setups (namely quantum networks and (single) qubit noise control) for which reinforcement learning (RL) agents manage to find interesting strategies. I hope to bridge this to the second “RL” (representation learning) in the title, where we ask about representations learned by AI for solving quantum tasks.

Automatic re-calibration of quantum devices by RL

Matias Bilkis
Computer vision center

During their operation, due to shifts in environmental conditions, devices undergo various forms of detuning from their optimal settings. Typically, this is addressed through control loops, which monitor variables and the device performance, to maintain settings at their optimal values. Quantum devices are particularly challenging since their functionality relies on precisely tuning their parameters. At the same time, the detailed modeling of the environmental behavior is often computationally unaffordable, while a direct measure of the parameters defining the system state is costly and introduces extra noise in the mechanism. In this study, we investigate the application of reinforcement learning techniques to develop a model-free control loop for continuous recalibration of quantum device parameters. Furthermore, we explore the advantages of incorporating minimal environmental noise models. As an example, the application to numerical simulations of a Kennedy receiver-based long-distance quantum communication protocol is presented.

Quantum Equilibrium Propagation for efficient training of quantum systems based on Onsager reciprocity

Clara Wanjura
Max Planck Institute for the Science of Light

The widespread adoption of machine learning and artificial intelligence in all branches of science and technology has created a need for energy-efficient, alternative hardware platforms. While such neuromorphic approaches have been proposed and realised for a wide range of platforms, physically extracting the gradients required for training remains challenging as generic approaches only exist in certain cases. Equilibrium propagation (EP) is such a procedure that has been introduced and applied to classical energy-based models which relax to an equilibrium. Here, I will show a direct connection between EP and Onsager reciprocity and exploit this to derive a quantum version of EP. This can be used to optimize loss functions that depend on the expectation values of observables of an arbitrary quantum system. Specifically, I will illustrate this new concept with supervised and unsupervised learning examples in which the input or the solvable task is of quantum mechanical nature, e.g., the recognition of quantum many-body ground states, quantum phase exploration, sensing and phase boundary exploration. We propose that in the future quantum EP may be used to solve tasks such as quantum phase discovery with a quantum simulator even for Hamiltonians which are numerically hard to simulate or even partially unknown. Our scheme is relevant for a variety of quantum simulation platforms such as ion chains, superconducting qubit arrays, neutral atom Rydberg tweezer arrays and strongly interacting atoms in optical lattices.

Autonomous tuning of spin qubits

Jonas Schuff
University of Oxford

For over two decades, semiconductor-based spin qubits have been a focus of quantum computing research. Despite significant progress, the complex manual tuning required to transform devices from their initial grounded state to fully operational qubits remains a substantial challenge. In this talk, I will present recent advancements in automating the tuning process of spin qubits. We have developed and tailored a range of machine learning techniques—including Bayesian optimization, deep learning, and computer vision—to address specific tasks within this process. I will discuss the problems we tackled, the solutions we implemented, and how our integrated framework enables grounded devices to become functional spin qubits without the need for human intervention.

Machine learning and ultracold quantum gases

Christof Weitenberg
Universität Hamburg, Institut für Quantenphysik

Ultracold quantum gases can benefit from modern machine learning tools in various ways, e.g., in the data evaluation or in the read-out process. As an example of the former, I will present our results on identifying phase transitions from experimental quantum-gas data enhanced via supervised and unsupervised machine learning. As an example of the latter, I will explain our recently introduced phase microscope and discuss possibilities for improving it via machine learning.

Improving robustness of quantum feedback control with reinforcement learning

Manuel Guatto
Forschungenszentrum Juelich

Obtaining reliable state preparation protocols is a key step towards practical implementation of many quantum technologies, and one of the main tasks in quantum control. In this work, different reinforcement learning approaches are used to derive a feedback law for state preparation of a desired state in a target system. In particular, we focus on the robustness of the obtained strategies with respect to different types and amount of noise. Comparing the results indicates that the learned controls are more robust to unmodeled perturbations with respect to simple feedback strategy based on optimized population transfer, and that training on simulated nominal model retain the same advantages displayed by controllers trained on real data. The possibility of effective off-line training of robust controllers promises significant advantages towards practical implementation.

Data-driven decoding of quantum error correcting codes using graph neural networks

Moritz Lange
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 decoders based on more efficient algorithms give sub-optimal performance. In addition, the accuracy will depend on the quality of models and estimates of error rates for idling qubits, gates, measurements, and resets, and will typically assume symmetric error channels. In this work, instead, we explore a model-free, data-driven, approach to decoding, using a graph neural network (GNN). The decoding problem is formulated as a graph classification task in which a set of stabilizer measurements is mapped to an annotated detector graph for which the neural network predicts 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, even if the matching decoder is given full information of the underlying error model. Although training is computationally demanding, inference is fast and scales approximately linearly with the space-time volume of the code. We also find that we can use large, but more limited, datasets of real experimental data [Google
Quantum AI, Nature [614, 676 (2023)]] for the repetition code, giving decoding accuracies that are on par with minimum weight perfect matching. The results show that a purely data-driven approach to decoding may be a viable future option for practical quantum error correction, which is competitive in terms of speed, accuracy, and versatility.

Towards a geodesic algorithm for optimizing variational circuits

Rafael Gomez Lurbe
University of Valencia

Currently, most optimization algorithms for updating the parameters of variational circuits rely on the gradient descent algorithm. Recently, a modification of this algorithm was proposed to include information about the geometry of the quantum state, known as Quantum Natural Gradient. This method involves using the Quantum Fisher Information Matrix, which can be obtained by taking the real part of the Quantum Geometric Tensor (QGT). Sometimes, the real part of the QGT is referred to as the Fubini-Study metric and serves as a metric in the projected Hilbert space. In our work, we propose using the Fubini-Study metric to follow a geodesic path (more efficient in time) during the optimization of a variational circuit. The result is an algorithm that alternates between gradient descent to choose the direction in which the cost is reduced, using this direction as initial conditions for the velocity in the geodesic part, and geodesic evolution, which reduces the optimization time. The algorithm was successfully tested on a toy model of one qubit to find the ground state of the Hamiltonian sigma-z.

Optimizing Local Hidden-Variable Models for Quantum Many-Body States

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, the understanding of non-locality remains limited, particularly in the context of measurements with an infinite number of possible settings (e.g., all projective measurements) and many-body systems. To address this, we developed a machine learning algorithm which optimizes local hidden-variable (LHV) models to reproduce the measurement statistics of many-body states. Our method efficiently produces LHV models for projective measurements of spin-1/2 systems, and provides estimates for the critical visibilities of two-qubit Werner and noisy three-qubit GHZ and W states. We find evidence suggesting that two-qubit subsystems in the ground states of translationally invariant Hamiltonians are local, while bigger subsystems are in general not. Our approach offers a tool for studying non-locality in all kinds of situations, such as ground states, thermal states, non-equilibrium states and states subject to different types of noise.

Phase Transitions in Clifford Circuits Through Reinforcement Learning

Giovanni Cemin
Max Planck Institute for the Physics of Complex Systems

In this research, we investigate the dynamics of entanglement in Clifford circuits by employing a reinforcement learning (RL) algorithm in competition with a random agent. The RL agent is designed to strategically place gates that decrease entanglement, while the random agent aims to increase entanglement. This interaction between the two agents results in a phase transition, the nature of which is influenced by the level of information accessible to the RL agent. By systematically varying the information provided to the RL agent, we analyze its impact on the phase transition characteristics. Our findings provide new insights into the interplay between entanglement manipulation and information constraints, shedding light on the fundamental mechanisms governing quantum circuit dynamics.

Deep Learning Time Dependent Density Functional Theory for Spin Hamiltonians

Emanuele Costa
University of Barcelona

Time-dependent density functional theory (TDDFT) is a powerful method for simulating the dynamics of quantum many-body systems by fully understanding the evolution of the one-particle density profile. Building on the generalization of TDDFT's main theorems to spin Hamiltonians, we investigate the implementation of the corresponding time-dependent Kohn-Sham equations. Utilizing the adiabatic approximation, we integrate a DL-functional with the Kohn-Sham equations to predict the dynamics of the transverse magnetization in a time-dependent transverse Ising Hamiltonian. Additionally, we address the non-adiabatic effects in the dynamics by employing an LSTM model to predict the exact Kohn-Sham driving.

Optimizing ZX-Diagrams with Deep Reinforcement Learning

Maximilian Nägele
Max Planck Institute for the Science of Light

ZX-diagrams are a powerful graphical language for the description of quantum processes with applications in fundamental quantum mechanics, quantum circuit optimization, tensor network simulation, and many more. The utility of ZX-diagrams relies on a set of local transformation rules that can be applied to them without changing the underlying quantum process they describe. These rules can be exploited to optimize the structure of ZX-diagrams for a range of applications. However, finding an optimal sequence of transformation rules is generally an open problem. In this work, we bring together ZX-diagrams with reinforcement learning, a machine learning technique designed to discover an optimal sequence of actions in a decision-making problem and show that a trained reinforcement learning agent can significantly outperform other optimization techniques like a greedy strategy, simulated annealing, and state-of-the-art hand-crafted algorithms. The use of graph neural networks to encode the policy of the agent enables generalization to diagrams much bigger than seen during the training phase.

Hamiltonian learning quantum magnets with non-local impurity tomography

Greta Lupi
Aalto University

Atomic-scale quantum magnets have gained significant attention due to their potential applications in quantum information processing and spintronics. Exploiting advancements in scanning tunneling microscopy (STM), it is now possible to measure and manipulate these systems with atomic precision. This project explores the application of supervised machine learning techniques to infer Hamiltonian parameters from STM data using non-local impurity tomography. We focus on two prototypical systems: a fermionic chain and a spin-½ chain, using the fermionic data as a benchmark to validate our approach for the spin chain. By training the NNs with theoretical datasets and subsequently fine-tuning them with noisy data, we aim to develop robust models capable of accurately predicting Hamiltonian parameters from experimental measurements. Our results demonstrate the efficacy of this approach in providing insights into the underlying interactions of quantum systems.

Automated Discovery of Coupled Mode Setups

Jonas Landgraf
Max Planck Institute for the Science of Light

In optics and photonics, a small number of building blocks, like resonators, waveguides, arbitrary 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 a time-consuming process where it remains uncertain whether the simplest possibility has indeed been found. In our work, we develop a discovery algorithm that automates this challenge. By optimizing the continuous and discrete system properties our auto- mated search identifies the minimal resources required to realize the requested scattering behavior. In the spirit of artificial scientific discovery, it produces a complete list of interpretable solutions and leads to generalizable insights, as we illustrate in several examples. This opens the door to- wards automated discovery of scattering setups for photonics, microwaves and optomechanics, with possible future extensions to periodic structures, sensing, and electronic devices.

Algorithms for quantum causal discovery

Jasleen Kaur
Macquarie University

Correlation does not imply causation—but then what does, especially in the quantum world? The discovery of causal relationships is an integral part of every scientific discipline. Even though causal modelling is a well developed field in classical physics, it is not straightforward in the quantum setting due to measurement playing a non-trivial role. Existing quantum causal discovery algorithms require knowledge about the complete process, which in turn requires process tomography to be performed with an informationally complete set of operations in order to detect the quantum causal model and the precise mechanisms behind the causal relations. The development of witnesses of quantum causal orders is an improvement over full process tomography, however it still requires a significant number of operations. In this work, we build an artificial neural network based machine learning algorithm that detects the causal structure of a multitime process using statistics from an informationally incomplete set of operations. The model exhibits exceptional accuracy, averaging to 98 percent and is robust against multinomial noise as well as Gaussian noise. Taking into account experimental and computational limitations, we further provide an alternative approach to detect causal connections given a quantum causal network of any size. The new approach is based on reduced processes and scales polynomially in comparison to exponential scaling in other approaches, thereby decreasing the computational complexity.

Machine Learning Density Functionals from Noisy Quantum Data

Felix Frohnert
Leiden University

The search for useful applications of noisy intermediate-scale quantum (NISQ) devices in quantum simulation has been hindered by intrinsic noise and the high costs associated with achieving high accuracy. A promising approach involves using quantum devices to generate training data for classical machine learning (ML) models. In this study, we explore the use of noisy quantum data in training an ML model for the Hubbard density functional, motivated by density functional theory. We benchmark various ML models against exact solutions, demonstrating that a neural-network ML model can successfully generalize and enhance small datasets subject to unbiased sampling noise, typical of NISQ algorithms. Conversely, when trained on data with algorithmic noise from the Variational Quantum Eigensolver, the model learns the noise bias. The trained models can be applied to solving new problem instances in a Kohn-Sham-like optimization scheme, benefiting from automatic differentiability and achieving reasonably accurate solutions on most problem instances. Our findings suggest a promising pathway for leveraging NISQ devices in practical quantum simulations, highlighting both the potential benefits and the challenges that need to be addressed for successful integration of quantum computing and ML techniques.

Learning disordered interactions between Rydberg atoms from experimental snapshots

Olivier Simard
École Polytechnique

Quantum simulators hold the potential to solve quantum many-body problems that are beyond the reach of classical computers, especially those that bear numerous degrees of freedom and feature long-range entanglement. To fulfill their prospects, quantum simulators must be fully controllable, allowing for precise tuning of the microscopic physical parameters that define their implementation. However, experimental control has intrinsic limits owing to imprecision of the optical tweezers and finite temperature, amongst other causes. One way to know accurately the Hamiltonian actually implemented is via Hamiltonian learning, where the goal is to develop a model that can infer the experimentally realized Hamiltonian in a scalable manner through measurements. In this work, we introduce a scalable approach to Hamiltonian learning using graph neural networks (GNNs). We employ the Density Matrix Renormalization Group (DMRG) method to generate ground-state snapshots of a quantum spin model for numerous sets of Hamiltonian parameters, whose correlation functions serve as input data to carry out the training. Our approach is exemplified by learning the Hamiltonian parameters of the quantum transverse-field Ising model defined on two-dimensional optical lattices. We train the GNNs on data from smaller lattice sizes (4x4, 5x5, and 6x6) and successfully predict Hamiltonian parameters for larger system sizes, extending up to 9x8 and 9x9, using experimental measurements such as local magnetization and nearest-neighbor spin correlation functions. This work shows that it could be possible to bridge the gap between numerical simulations and experiments with Rydberg atoms in optical tweezers, as well as allow fast real-time adjustment of the optical tweezers to implement the targeted Hamiltonian.

Deep RL Algorithms for the Average Reward Objective

Jacob Adamczyk
University of Massachusetts Boston

In reinforcement learning (RL), the discounted objective is commonly used for its corresponding popular algorithms and guaranteed convergence. However, in physical systems where the non-linear discounting of rewards or energy is often unprincipled, the average reward formulation may be more appropriate. Despite its relevance, average reward RL has received less attention, particularly in the deep RL setting, where few algorithms exist. We introduce new value-based algorithms for entropy-regularized average reward RL and demonstrate their effectiveness in large-scale problems. We also present an algorithm relevant for quantum ground state problems based on a novel "annealing" strategy. By connecting with ideas from statistical mechanics such as free energy and quasi-stationary distributions, we offer fresh insights into the intersection of RL and physics.

Quantum circuit synthesis with diffusion models

Florian Fürrutter
University of Innsbruck

Achieving the expected advantages of quantum computing relies on efficiently translating quantum operations into viable physical realizations. In this work, we use state-of-the-art generative machine learning models, specifically denoising diffusion models (DMs), to facilitate this transformation. Utilizing text-conditioning, we steer the model to produce desired quantum operations within gate-based quantum circuits. Notably, during training, DMs allow sidestepping the exponential overhead inherent in the classical simulation of quantum dynamics — a consistent bottleneck in preceding machine learning 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.

Machine learning-aided steering certification

Arthur Dutra de Oliveira
UNICAMP

Quantum steering is an essential resource in various quantum technologies applications, yet determining whether a state is steerable remains a difficult challenge. Although machine learning techniques have been explored for this purpose, they typically provide probabilistic answers that, even when highly accurate, lack certainty. This probabilistic nature limits their applicability in scenarios where steerability certification is required, such as in security proofs or computational tasks. In this work, we introduce a novel tool that combines state-of-the-art steering detection algorithms with a machine learning model, achieving speed improvements without sacrificing the deterministic nature of the certification process. Our approach is specifically optimized for two-qubit states and demonstrates exceptional utility in scenarios requiring the certification of large numbers of states, such as volume estimations, brute force searches in state space, or decentralized computing applications, paving the way for more efficient and reliable quantum state certification methods.

Reinforcement Learning Based Quantum Circuit Optimization via ZX-Calculus

Jan Nogué Gómez
Qilimanjaro Quantum Tech

We propose a novel Reinforcement Learning (RL) method for optimizing quantum circuits using graph-theoretic simplification rules of ZX-diagrams. The agent, trained using the Proximal Policy Optimization (PPO) algorithm, employs Graph Neural Networks to approximate the policy and value functions. We demonstrate the capacity of our approach by comparing it against the best performing ZX-Calculus-based algorithm for the problem in hand. After training on small Clifford+T circuits of 5-qubits and few tenths of gates, the agent consistently improves the state-of-the-art for this type of circuits, for at least up to 80-qubit and 2100 gates, whilst remaining competitive in terms of computational performance. Additionally, we illustrate its versatility by targeting both total and two-qubit gate count reduction, conveying the potential of tailoring its reward function to the specific characteristics of each hardware backend. Our approach has been succesfully used to optimize a circuit of 12 qubits for Relativistic VQE calculations of molecular electric dipole moments on trapped ion quantum hardware (IonQ).

Preparing Schrödinger cat states using a neural network

Pavlo Bilous
Max Planck Institute for the Science of Light

We show how a neural network (NN) can be used to tailor control pulses for preparation of quantum states in a cavity coupled to a qubit. Instead of considering each quantum state separately, the approach focuses on continuous families of states. Once trained on states randomly selected from the family, the NN is able to generate correct control pulses for any state from the family without further optimization or retraining. As a concrete demonstration of the method, we applied it to Schrödinger cat states and performed an experimental verification of the result.

Finding optimal experimental parameters for superconducting qubit readout

Andras Di Giovanni
Karlsruher Institut für Technologie

Scaling up few-qubit superconducting qubit devices is proving to be a challenging problem in current quantum computing efforts. Multiplexed readout aims to overcome one of the problems of scaling up, reducing the overall circuit footprint as well as the number of required signal lines and amplifier chains. This is achieved by coupling multiple readout resonators to a mutual transmission line, enabling simultaneous readout of the multi-qubit system. Although very widespread, this method of readout can lead to significant crosstalk particularly in the case of resonators being only weakly detuned in frequency. A second problem is the nontrivial dependence of amplification on the frequency of the resonators. To study these problems, we implement a detector tomography based method for finding optimal readout parameters that have a low correlation, while not significantly affecting the quality of single shot readout. This approach is useful as it also gives the direction on the correlation and takes into account the full positive operator-valued measure of the system. We employ Bayesian optimization, which minimizes a cost function containing the relevant experimental properties of a good measurement. We implement the method on a superconducting qubit chip with multiple Xmon qubits and test its ability to find optimal operating parameters for two resonators that are very close in frequency.

Adaptive Quantum Sensing in Nitrogen-Vacancy Centres in Diamond at Room Temperature

Stefan Todd
Heriot-Watt University

Quantum sensors utilising single electron spins have opened unprecedented opportunities to detect and measure nanoscale phenomena in condensed matter physics, telecommunication, electronics, biomedical research [1-4]. These measurements however necessitate long timescales compared to ensemble approaches and this problem is exacerbated at room temperature, where single-shot spin readout is currently not available. The duration of these experiments is therefore a limiting factor in applications, which can be circumvented by online techniques that allow us to cut out 'insignificant' results, speeding data collection up. I will describe adaptive methods (based on Bayesian inference, reinforcement learning and neural networks) to optimise quantum sensors to seek out regions of impactful settings, in real-time. We have developed an experimental adaptive setup, based on the single electron spin associated to a nitrogen-vacancy (NV) centre in diamond, that uses Bayesian inference to learn a given parameter and optimise settings on the fly [5]. In particular, I will describe methodologies to improve the estimation of static (DC) magnetic fields through adaptive Ramsey pulse sequences, which allows us to measure the Larmor precession frequency of the spin sensors (proportional to the applied magnetic field). I will discuss our experiments in optimising the real-time choice of experimental settings based on reinforcement learning [6] and to include experimental imperfections by learning noise with a neural network, in a graybox approach [7].

[1] - Ho, K. O., et al. (2022). Functional Diamond
[2] - Dréau, A., et al. (2018). Physical Review Applied
[3] - Kim, D., et al. (2019). Nature Electronics
[4] - Schirhagl, R., et al. (2014). Annual Review of Physical Chemistry
[5] - Arshad, M. J., et al. (2024). Physical Review Applied
[6] - Belliardo, F., et al. (2022). Physical Review Research
[7] - Youssry, A., et al. (2024). Npj Quantum Information

Reservoir Computing for Time Series Prediction

Thomas Cope
IQM

In this work, we present recent empirical results on predicting time series data using quantum circuits as a reservoir. We discuss the structure of circuits used to do this; and the performance both simulated and on real hardware.

Machine Learning in Loop Quantum Gravity

Waleed Sherif
FAU Erlangen-Nürnberg

In this talk, I will introduce the first case of machine learning techniques being used to address the arguably most difficult and open challenge in loop quantum gravity: finding and interpreting solutions to the quantum Hamiltonian constraint. While this problem has proven resistant to both analytical and traditional numerical methods due to its inherent complexity, I will demonstrate that the application of neural network quantum states (NQS) in this context opens a possibility to tackle this problem. I will do so by showcasing an proof-of-concept example whereby the NQS ansatz is used to solve the constraints of loop quantised 3-dimensional Euclidean gravity in a certain weak coupling limit, arriving at solutions for complicated constraints which lie in extremely large spaces both efficiently and accurately. Additionally, I will touch on ongoing work and future directions for integrating neural networks into quantum gravity research.

Zero-Temperature Neural Monte Carlo simulations of quantum spin glasses

Luca Brodoloni
University of Camerino

Simulating the low-temperature properties of disordered quantum Ising models is a paradigmatic problem in condensed matter physics, and it has recently gained strong interest in the context of quantum-enhanced optimization performed via quantum annealers. In this study, we employ a continuous-time projection quantum Monte Carlo algorithm enhanced by self-learning neural network-based guiding wave functions to simulate the ground state properties of quantum spin glasses, focusing on both mean-field and short-range models. Specifically, we explore the two-dimensional Edwards-Anderson Hamiltonian with a transverse field and Gaussian nearest-neighbor couplings. Our approach effectively mitigates population control bias and provides unbiased estimates of ground-state properties, enabling the investigation of regimes and system sizes previously inaccessible. Through a finite-size scaling analysis, we estimate the critical transverse field and the critical exponents associated with the spin-glass transition. For the latter two remarkable agreement is found with recent estimates from the literature for different random couplings . Additionally, we identify the quantum critical point of the Edwards-Anderson model and investigate the spin-overlap distribution within the spin-glass phase, revealing a non-trivial double-peak structure indicative of Replica Symmetry Breaking (RSB) at $T = 0$.

Unitary Synthesis of Clifford+T Circuits with Reinforcement Learning

Christian Ufrecht
Fraunhofer IIS

We present a deep reinforcement learning approach for synthesizing unitaries into quantum circuits. Unitary synthesis aims to identify a quantum circuit that represents a given unitary while minimizing circuit depth, total gate count, a specific gate count, or a combination of these factors. While past research has focused predominantly on continuous gate sets, synthesizing unitaries from the parameter-free Clifford+T gate set remains a challenge. Although the time complexity of this task will inevitably remain exponential in the number of qubits for general unitaries, reducing the runtime for simple problem instances still poses a significant challenge. We apply the tree-search method Gumbel AlphaZero to solve the problem for a subset of exactly synthesizable Clifford+T unitaries. Our method effectively synthesizes circuits for up to five qubits generated from randomized circuits with up to 60 gates, outperforming existing tools like QuantumCircuitOpt and MIN-T-SYNTH in terms of synthesis time for larger qubit counts. Furthermore, it surpasses Synthetiq in successfully synthesizing random, exactly synthesizable unitaries. These results establish a strong baseline for future unitary synthesis algorithms.

Quasiparticle Dispersions in SU(N) Models with Neural Quantum States

Annika Böhler
LMU Munich

The central challenge in numerical simulations of quantum many-body systems is posed by the exponential growth of the Hilbert space dimension with system size. Recently, neural quantum states (NQS) have emerged as a promising new variational class, where the many- body amplitude is parametrized in terms of an artificial neural network. Motivated by the crucial role of symmetries in many-body physics, we aim to enhance the expressive capacity of NQS by utilising group convolutional neural networks, a highly symmetric architecture that can be understood as a generalisation of standard convolutional neural networks. Drawing on concepts from representation theory, we use this architecture to carry out a ground state search in a fixed symmetry sector. We apply this architecture to find the dispersion relation of SU(2) and SU(3) tJ-models.

AI-powered discovery and design of novel quantum experiments

Carlos Ruíz-González
Max Planck Institute for the Science of Light

Photonic technologies provide advanced sensors, secure communications, and quantum-enhanced computation. Some of these endeavors require generating specific quantum states or efficiently performing quantum tasks. The design of such setups, historically powered by human creativity, is being slowly automated by novel algorithms. Unfortunately, these tools are often restricted to very specific use cases and are difficult to generalize, which limits their practical implementation. To overcome these challenges, we developed PyTheus, a highly-efficient, open-source digital discovery framework based on a graph-based representation of optical designs, including a wide range of modern experimental devices. PyTheus produces interpretable results to solve complex experimental problems, like entanglement generation and distribution. Aiming for the simplest solutions, our software provides deeper understanding to human researchers to generalize their findings. This makes PyTheus a powerful tool for the development of quantum optics and related technologies.

Large-Scale Discovery of Experimental Designs in Super-Resolution Microscopy with XLuminA

Carla Rodríguez
Max Planck Institute for the Science of Light

Driven by human ingenuity and creativity, the discovery of super-resolution techniques, which circumvent the classical diffraction limit of light, represent a leap in optical microscopy. However, the vast space encompassing all possible experimental configurations suggests that some powerful concepts and techniques might have not been discovered yet, and might never be with a human-driven direct design approach. Thus, AI-based exploration techniques could provide enormous benefit, by exploring this space in a fast, unbiased way. We introduce XLuminA, an open-source computational framework developed using JAX, which offers enhanced computational speed enabled by its accelerated linear algebra compiler (XLA), just-in-time compilation, and its seamlessly integrated automatic vectorization, auto-differentiation capabilities and GPU compatibility. Remarkably, XLuminA demonstrates a speed-up of 4 orders of magnitude compared to well-established numerical optimization methods. We showcase XLuminA's potential by rediscovering two foundational techniques in advanced microscopy, together with new superior experimental layouts. Ultimately, XLuminA identified a novel experimental blueprint featuring sub-diffraction imaging capabilities. This work constitutes and important step in AI-driven scientific discovery of new concepts in optics and advanced microscopy.

Hybrid discrete-continuous compilation of trapped-ion quantum circuits with deep reinforcement learning

Francesco Preti
Forschungszentrum Juelich

Shortening quantum circuits is crucial to reducing the destructive effect of environmental decoherence and enabling useful algorithms. We demonstrate an improvement in such compilation tasks via a combination of using hybrid discrete-continuous optimization across a continuous gate set, and architecture-tailored implementation. The continuous parameters are discovered with a gradient-based optimization algorithm, while in tandem the optimal gate orderings are learned via a deep reinforcement learning algorithm, based on projective simulation. To test this approach, we introduce a framework to simulate collective gates in trapped-ion systems efficiently on a classical device. The algorithm proves able to significantly reduce the size of relevant quantum circuits for trapped-ion computing. Furthermore, we show that our framework can also be applied to an experimental setup whose goal is to reproduce an unknown unitary process.

Noise classification in two qubit system with Machine learning

Shreyasi Mukherjee
University of Catania

We investigate a machine learning based classification of noise acting on a two qubit system with the aim of detecting noise correlations, and the interplay with Markovianity. We study a ultra strongly coupled two qubit system which can be reduced to an effective three-level system under some particular choice of parameters and approximations. We then control the system using the well known Stimulated Raman Adiabatic Passage (STIRAP) protocol to achieve effective population transfer from ground state to the doubly excited state of the system. We use the efficiency of population transfer under different combinations of external drive amplitudes as inputs to train a feedforward neural network. We show that supervised learning can classify different types of classical noise acting locally on the qubits and affecting the overall system dynamics. In particular as noise models we take into account three non-Markovian (quasi-static correlated, anti-correlated and uncorrelated) and two Markovian noises (correlated and anti-correlated) that are classified with significant accuracy. Our approach is robust to statistical measurement errors and retains its effectiveness for physical measurements where only a limited number of samples is available.

Boosted quantum teleportation

Simone DAurelio
FMQ Uni Stuttgart

Quantum Teleportation is a building block for Quantum Information tasks. Even with its simplicity, the ability to perform Quantum Teleportation has implication in many more applications, ranging from quantum communication to quantum computation. While photons are good candidates for both tasks, performing teleportation With linear optics present limitations when it comes to projecting qubits into the Bell states basis. In fact, the efficiency of Bell-State measurements is limited at 50% with just linear optics. Here we present the basis of a Boosted Bell-state measurement and its application in a Quantum Teleportation protocol, explaining the details of the experimental setup. With the use of SPDC sources emitting degenerate photons at telecom-C wavelength, we teleport the polarization state of a photon, proving fidelities between the teleported states and the expected ones higher than the classical fidelity bound.

Structure motifs in interpretable quantum algorithm design

Davide Bincoletto
University of Augsburg

One of the most studied algorithms for modern quantum computers in chemistry applications is the Variational Quantum Eigensolver (VQE). This hybrid quantum-classical algorithm combines state preparation on a quantum computer with optimization on a classical computer to efficiently approximate the ground state energy of a quantum system. State preparation involves constructing an ansatz, a series of parameterized quantum gates, to represent the system's wave function. A crucial aspect is designing ansätze that are accurate, resource-efficient, and interpretable. In this context, an important heuristic is the Graph-Based Circuit Design. It encodes the system chemical information in a series of graph structures, where vertices represent atoms and edges represent interactions between them. These structures are inspired by Valence Bond Theory, i.e. each vertex has one edge in order to produce only paired interactions between atoms. Each of these graphs is then translated into parametrized quantum gates of topology: basis change - double excitation - backward basis change. This heuristic satisfies the three requirements stated before, especially due to local connectivity. This work introduces a method to analyze the orbital space transformation induced by the basis change process, aiding in the ansatz interpretability. This is given by leveraging the matrix representation of a subspace rotation operator, also called Givens rotation. In addition, some possible extensions of this ansatz are presented in order to show its flexibility and accuracy. These are obtained using quantum gates which produce more general basis change and electrons double excitations compared to the graph structures.

Routing quantum circuits with AlphaZero deep exploration

Marvin Richter
Chalmers University of Technology

Compiling a quantum circuit for specific quantum hardware presents a significant challenge due to the low connectivity of physical qubits and limited coherence time in current quantum processing units. Qubit routing, the circuit transformation from the original circuit to a connectivity-adapted version, employs SWAP gates to meet two-qubit gate requirements. This process must minimize depth overhead to maximize the utilization of limited resources while enabling the execution of reasonably sized quantum circuits on target hardware. The large, combinatorial search space for such circuit transformations, coupled with a high branching factor, often leads existing algorithms for higher qubit numbers to conduct only superficial searches, resulting in suboptimal solutions. Our method combines two frameworks to address this challenge. First, a transformer network trained via Monte Carlo Tree Search (MCTS) proposes mappings at each layer of the quantum circuit. Second, a routing algorithm based on graph matchings realizes these proposed mappings. This hybrid approach ensures strategic mappings while guaranteeing consistent and deterministic routing through an algorithm that finds optimal routing sequences for a grid architecture. For arbitrary qubit topologies, the transformer network can directly place SWAP gates instead of changing the mappings. The MCTS, guided by a transformer neural network, allows for deeper searches and potentially better solutions than existing heuristic algorithms, such as SABRE. We present benchmarks comparing our method to other algorithms, demonstrating its performance in finding shorter, more efficient solutions.

Computational battery modeling with differentiable quantum circuits

Nicolò Toscano
Pasqal

Many scientific problems are formulated in the form of Partial Differential Equations (PDEs). The most popular approach to solve these systems are mesh-based methods, which are powerful and reliable but with severe scaling issues when complexity increases. Physics-Informed NNs (PINNs) are part of the emerging field of Scientific Machine Learning (SciML) and are proposed as an alternative paradigm, using a NN as universal function approximator (UFA) to model the solution of the target PDE. The key advantages of this approach are its flexibility, the opportunity to directly embed system properties as well as use experimental data for training or regularization and, finally, its inference capabilities. Our research showcases the use of Differentiable Quantum Circuits (DQC), a hybrid quantum-classical algorithm inspired from PINNs, applied to solving systems of PDEs relevant to battery simulation. We validate the algorithm's accuracy in generating solutions for multiple variables simultaneously over several connected subdomains. Given the high expressivity provided by its large basis-set (which scales exponentially with the register size) and the low depth required to converge, this algorithm can perform well on complex tasks that are challenging for its classical counterparts while remaining NISQ-compatible.

Quantum Process Tomography of Structured Optical Gates using Convolutional Neural Networks

Tareq Jaouni
University of Ottawa

The characterization of a unitary gate is experimentally accomplished via Quantum Process Tomography, which combines the outcomes of different projective measurements to reconstruct the underlying operator. The process matrix is typically extracted from maximum-likelihood estimation. Recently, optimization strategies based on evolutionary and machine-learning techniques have been proposed. Here, we investigate a deep-learning approach that allows for fast and accurate reconstructions of space-dependent SU(2) operators, only processing a minimal set of measurements. We train a convolutional neural network based on a scalable U-Net architecture to process entire experimental images in parallel. Synthetic processes are reconstructed with average fidelity above 90%. The performance of our routine is experimentally validated on complex polarization transformations. Our approach further expands the toolbox of data-driven approaches to Quantum Process Tomography and shows promise in the real-time characterization of complex optical gates.

Training hybrid quantum-classical neural networks for the recognition of topological phases

Markus Hoffman
FAU Erlangen-Nürnberg

With the increasing complexity of quantum computers, standard methods for the characterization of quantum states based on direct measurement and classical post-processing have become impractical due to large measurement costs. Quantum neural networks based on parametrized quantum circuits are a promising tool for characterizing quantum states with reduced sample complexity. However, their application has been limited to only proof-of-principle demonstrations as current noisy intermediate-scale quantum devices can execute only short quantum circuits. To overcome this challenge, we introduce hybrid quantum-classical neural networks that consist of a parametrized quantum circuit, measurement and classical post-processing. Using supervised and unsupervised learning, we train these hybrid neural networks to recognize the topological phase of the nine-qubit surface code in an external field. We show that the trained parametrized quantum circuit maximizes the Fisher information and Jensen-Shannon divergence of the measurement outcome distributions. Subsequently, we determine the Bayes-optimal prediction of quantum phases in classical post-processing. We demonstrate that, using unsupervised learning by confusion, a more precise indicator of the phase boundary is obtained compared to supervised learning. Since these hybrid neural networks require only short-depth quantum circuits, they open the way for the efficient characterization of quantum states on current and near-term quantum computers.

Neural quantum states for operators

Marjan Macek
FAU Erlangen-Nürnberg

A considerate development in recent years in variational Monte Carlo has been the application of neural networks as the wave-function ansatz. There are several neural quantum states (NQS) ansätze for wave functions in closed systems, but NQS ansätze for density matrices are lacking either in expressiveness or in capturing the properties of the density matrices. Typically, NQS that lack the properties of the target exhibit a lower performance than those that have them. Furthermore, a NQS, which is not guaranteed to be Hermitian and positive definite, may even describe nonphysical states. In a steady-state calculation, we are guaranteed to converge towards the physical steady-state. However, this is not the case during the time evolution. Positive definiteness is the most challenging for ansätze. We propose doing a time evolution in the Heisenberg picture, where we time evolve the operator. We will present results for a transverse-field Hamiltonian with dissipation.

How non-classical is a given quantum state?

Martina Jung
Friedrich-Schiller-Universität Jena

Non-classicality, defined and understood in the sense of quantum optics, must be seen as a resource: If a non-classical state – characterized by a partially negative Glauber-Sudarshan P-representation – is mixed with vacuum in a beamsplitter, the resulting state will be entangled. Hence, quantifying the non-classicality of a quantum state is crucial to gauge its potential for quantum advantage in an experiment, for instance in a Boson Sampler. From a quantum information perspective, the non-classicality of a state is given by its distance to the set of classical states. Based on this geometric notion of non-classicality, our variational approach minimizes over the coarse-grained set of classical states which is parametrized by a neural network. The network is then trained to approximate the closest classical state to a given target state characterized by its phase-space representation. Recognizing that this measure strongly depends on the chosen distance measure, we also implement a data-driven approach which quantifies the non-classicality of a state by the ability of a neural network to distinguish the state from a classical one. In this approach, samples from photon-number measurements are arranged to form a two-dimensional image that is processed by a classifier-Convolutional Neural Network (CNN). The question of how non-classical a quantum state is, is then rephrased as "How many samples does the model have to see to recognize the state as being non-classical?"

Lottery ticket neural network quantum states

Brandon Barton
ETH Zurich

Despite significant innovations in the development of neural network quantum states (NQS), the appropriate choice of ansatz, interpretability of parameters, and scaling behaviors of NQS remain unclear in many cases. This work applies neural network compression techniques to investigate the capabilities and limitations of NQS, drawing inspiration from the Lottery Ticket Hypothesis. We demonstrate that unstructured iterative magnitude pruning with weight rewinding (IMP-WR) is an effective NQS compression technique for ground state problems in 2D quantum spin models, including the transverse-field Ising and Toric code models. Using fully-connected and convolutional neural networks trained with variational Monte Carlo sampling, we identify sparse sub-networks that match or exceed the performance of their dense counterparts while retaining less than 15% of the original parameters. Our methods reveal highly interpretable features corresponding to distinct quantum phases of matter through careful monitoring of network parameters during pruning. This work establishes a framework for discovering interpretable NQS architectures, achieving significant parameter reduction while maintaining or enhancing model performance.

Gradient-ascent pulse engineering with feedback

Vittorio Peano
Max Planck Institute for the Science of Light

Efficient approaches to quantum control and feedback are essential for quantum technologies, from sensing to quantum computation. Open-loop control tasks have been successfully solved using optimization techniques, including methods such as gradient-ascent pulse engineering (GRAPE) , relying on a differentiable model of the quantum dynamics. For feedback tasks, such methods are not directly applicable, since the aim is to discover strategies conditioned on measurement outcomes. In this work, we introduce feedback GRAPE, which borrows some concepts from model-free reinforcement learning to incorporate the response to strong stochastic (discrete or continuous) measurements, while still performing direct gradient ascent through the quantum dynamics. We illustrate its power by considering various scenarios based on cavity-QED setups. Our method yields interpretable feedback strategies for state preparation and stabilization in the presence of noise. Our approach could be employed for discovering strategies in a wide range of feedback tasks, from calibration of multiqubit devices to linear-optics quantum computation strategies, quantum-enhanced sensing with adaptive measurements, and quantum error correction.

Quantum Wasserstein Compilation: Unitary Compilation using the Quantum Earth Mover's Distance

Abhishek Dubey
Fraunhofer IIS

Despite advances in the development of quantum computers, the practical application of quantum algorithms remains outside the current range of so-called noisy intermediate-scale quantum devices. Now and beyond, quantum circuit compilation (QCC) is a crucial component of any quantum algorithm execution. Besides translating a circuit into hardware-specific gates, it can optimize circuit depth and adapt to noise. Variational quantum circuit compilation (VQCC) optimizes the parameters of an ansatz according to the goal of reproducing a given unitary transformation. In this work, we present a VQCC-objective function called the quantum Wasserstein compilation (QWC) cost function based on the quantum Wasserstein distance of order 1. We show that the QWC cost function is upper bound by the average infidelity of two circuits. An estimation method based on measurements of local Pauli-observable is utilized in a generative adversarial network to learn a given quantum circuit. We demonstrate the efficacy of the QWC cost function by compiling a single-layer hardware efficient ansatz (HEA) as both the target and the ansatz and comparing other cost functions such as the Loschmidt echo test (LET) and the Hilbert-Schmidt test (HST). Finally, our experiments demonstrate that QWC as a cost function can mitigate the barren plateaus for the particular problem we consider.

Gradient-ascent pulse engineering with feedback

Vittorio Peano
Max Planck Institute for the Science of Light

Efficient approaches to quantum control and feedback are essential for quantum technologies, from sensing to quantum computation. Open-loop control tasks have been successfully solved using optimization techniques, including methods such as gradient-ascent pulse engineering (GRAPE) , relying on a differentiable model of the quantum dynamics. For feedback tasks, such methods are not directly applicable, since the aim is to discover strategies conditioned on measurement outcomes. In this work, we introduce feedback GRAPE, which borrows some concepts from model-free reinforcement learning to incorporate the response to strong stochastic (discrete or continuous) measurements, while still performing direct gradient ascent through the quantum dynamics. We illustrate its power by considering various scenarios based on cavity-QED setups. Our method yields interpretable feedback strategies for state preparation and stabilization in the presence of noise. Our approach could be employed for discovering strategies in a wide range of feedback tasks, from calibration of multiqubit devices to linear-optics quantum computation strategies, quantum-enhanced sensing with adaptive measurements, and quantum error correction.

Employing tensor networks and Riemannian quantum circuit optimization for fermionic Hamiltonian simulation

Isabel Nha Minh Le
Technical University of Munich

Simulating fermionic systems is relevant for many fields like quantum chemistry or condensed matter physics. Trotterization is a simple and common tool to approximate the time evolution on a quantum computer. However, accurate simulations for long times are restricted by the resulting deep quantum circuits. The goal of this work is to increase accuracy under constant depth, utilizing only classical resources before employing any quantum device for simulation. We start from a fermionic swap network that implements a Trotter step using a brickwall circuit layout. To further increase the accuracy of the initial circuit, Riemannian optimization is employed to optimize the two-qubit gates under unitary constraints. The reference time evolution operator can efficiently be expressed as a matrix product operator for short enough times t. By interpreting the quantum circuit as a tensor network, a suitably chosen cost function and the corresponding gradient can be evaluated using tensor network methods. Executing the optimized circuit repetitively on a quantum device enables simulation times T>>t. We apply our method to molecular Hamiltonians and the 1D Fermi-Hubbard for 60 spin orbitals and t=0.3.

The information theory of machine learning for decoding QECCs

Evan Peters
University of Waterloo

Quantum Error-Correcting Codes (QECCs) are necessary for doing useful tasks with quantum computers, but optimally decoding QECCs is computationally hard. Meanwhile, Machine Learning (ML) has produced good heuristic solutions for a variety of hard problems, and has therefore received growing attention as a tool for decoding QECCs. But what, specifically, does it mean to do ML for decoding QECCs? In this work, we review how decoding QECCs in the stabilizer picture is described by a classical hypothesis testing problem. We survey the literature on ML for decoding QECCs, showing how the many existing approaches to this problem fit nicely into a few categories using our formalism. We will then discuss the ultimate limits of ML for decoding QECCs. We investigate upper and lower bounds for the accuracy of any such scheme using single-shot information theory, with a specific focus on practical applications - for example, by lower bounding how much harder it is to decode QECCs as your quantum computer gets hotter.

Quantum Convolutional Neural Network for Phase Recognition in Two Dimensions

Leon Sander
FAU Erlangen-Nürnberg

Quantum convolutional neural networks (QCNNs) are quantum circuits for recognizing quantum phases of matter at low sampling cost and have been designed for condensed matter systems in one dimension. Here we construct a QCNN that can perform phase recognition in two dimensions and correctly identify the phase transition from a Toric Code phase with $\mathbb{Z}2$-topological order to the paramagnetic phase. The network also exhibits a noise threshold up to which the topological order is recognized. Our work generalizes phase recognition with QCNNs to higher spatial dimensions and intrinsic topological order.

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's unclear how finding general design concepts can be automated. Here, we addresses 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. To demonstrate the effectiveness of our approach, we uncover previously unknown experimental generalizations of important quantum states, e.g. from condensed matter physics. The underlying methodology of meta-design can naturally be extended to fields such as materials science or engineering.

Feature Design for Learning-based Quantum Error Mitigation

Vincenzo Lipardi
Maastricht University

Quantum error mitigation is crucial for improving the reliability of quantum computations on near-term quantum devices. On of the main challenges in this field is reducing the sampling overhead, which refers to the large number of measurements required to accurately mitigate quantum noise. Previous works have shown the potential of machine learning models that trade off training time with inference time, offering promising alternatives for error mitigation. In a recent work by Liao H. et al., the authors achieved promising experimental results on both unstructured random circuits and structured circuits, by providing a comparison between several machine learning models and the Zero-Noise Extrapolation method. Random Forest turned out to be the most performant even compared with Graph Neural Networks, in terms of both computation time and performance. As suggested by the authors in the future directions, delving into the study of the quantum circuit features in relation to the model may provide insights for designing more performant solutions. We will focus on the design of quantum circuit features with classical shadows, a technique that efficiently captures quantum state properties in a classical representation. Moreover, we will investigate hardware-related features, such as pulse-level information in superconductive quantum hardware. By exploring various feature sets we aim to identify those that best perform in relation to different machine learning model in order to further investigate the potential of classical machine learning models for quantum error mitigation.

Neural Network Quantum States for the Interacting Hofstadter Model with Higher Local Occupations and Long-Range Interactions

Fabian Döschl
LMU Munich

Due to their immense representative power, neural network quantum states (NQS) have gained significant interest in current research. In recent advances in the field of NQS, it has been demonstrated that this approach can compete with state-of-the-art numerical techniques, making NQS a compelling alternative, in particular for the simulation of large, two-dimensional quantum systems. In this study, we show that recurrent neural network (RNN) wave functions can be employed to study systems relevant to current research in quantum many-body physics. Specifically, we employ a 2D tensorized gated RNN to explore the bosonic Hofstadter model with a variable local Hilbert space cut-off and long-range interactions. At first, we benchmark the RNN-NQS for the Hofstadter-Bose-Hubbard (HBH) Hamiltonian on a square lattice. We find that this method is, despite the complexity of the wave function, capable of efficiently identifying and representing most ground state properties. Afterwards, we apply the method to an even more challenging model for current methods, namely the Hofstadter model with long-range interactions. This model describes Rydberg-dressed atoms on a lattice subject to a synthetic magnetic field. We study systems of size up to 12×12 sites and identify three different regimes by tuning the interaction range and the filling fraction ν. In addition to phases known from the HBH model at short-ranged interaction, we observe bubble crystals and Wigner crystals for long-ranged interactions. Especially interesting is the evidence of a bubble crystal phase on a lattice, as this gives experiments a starting point for the search of clustered liquid phases, possibly hosting non-Abelian anyon excitations. In our work we show that NQS are an efficient and reliable simulation method for quantum systems, which are the subject of current research. In particular, we demonstrate the ability of this method to simulate challenging systems with long-range interactions.

Interesting Scientific Idea Generation using knowledge graphs and large language models

Xuemei Gu
Max Planck Institute for the Science of Light

The rapid growth of scientific literature makes it challenging for researchers to identify novel and impactful ideas, especially across disciplines. Modern artificial intelligence (AI) systems offer new approaches, potentially inspiring ideas not conceived by humans alone. But how compelling are these AI-generated ideas, and how can we improve their quality? Here, we introduce SciMuse, which uses 58 million research papers and a large-language model to generate research ideas. We conduct a large-scale evaluation in which over 100 research group leaders - from natural sciences to humanities - ranked more than 4,400 personalized ideas based on their interest. This data allows us to predict research interest using (1) supervised neural networks trained on human evaluations, and (2) unsupervised zero-shot ranking with large-language models. Our results demonstrate how future systems can help generating compelling research ideas and foster unforeseen interdisciplinary collaborations.