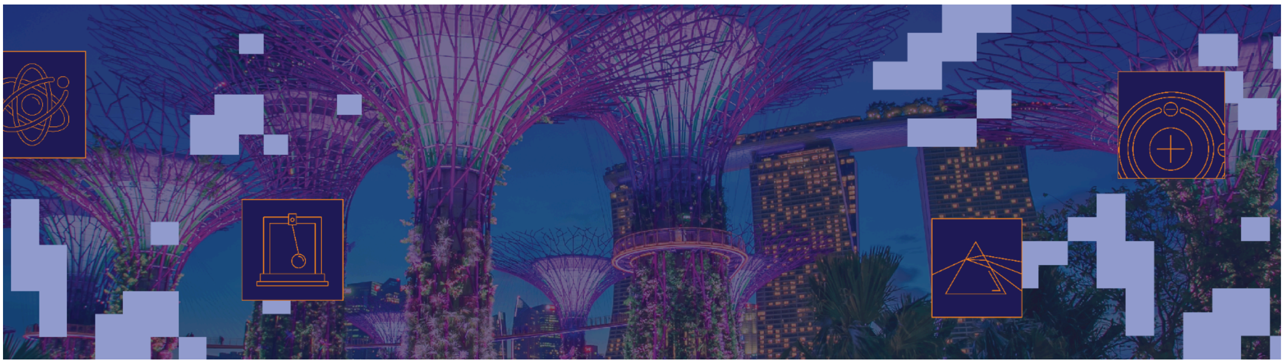

2025

SINGAPORE

CONFERENCE PROGRAMME

9th International Conference on Quantum Techniques in Machine Learning

Singapore, 16-21 November 2025

Quantum Techniques in Machine Learning (QTML) is a leading international conference at the forefront of quantum science and machine learning. Held annually, it brings together researchers and industry experts to explore how quantum computing can transform learning, optimization, and data-driven discovery. Through a series of scientific talks and discussions, QTML fosters collaboration and advances research on the interplay between quantum mechanics and machine learning, from foundational theory to real-world applications.

This programme booklet collates details of keynote, invited, and submitted talks, including those chosen by the programme committee as highlights, tutorials and posters.

Schedule

Please refer to the QTML 2025 website: https://qtml2025.cqt.sg/schedule/

Table of Contents

QTML Keynote Talks

Title: Quantum Probe Tomography by Sitan Chen, Harvard University

Abstract: Characterizing quantum many-body systems is a fundamental problem across physics, chemistry, and materials science. While significant progress has been made, many existing Hamiltonian learning protocols demand digital quantum control over the entire system, creating a disconnect from many real-world settings. Can one learn the parameters of a many-body Hamiltonian using a single local probe access to a small subsystem of a many-body thermal state undergoing time evolution? I will describe a new combination of tools from algebraic geometry and smoothed analysis that yields a provably correct algorithm for learning generic Hamiltonians in various physically natural families even in this setting. This demonstrates that robust Hamiltonian learning remains achievable even under severely constrained experimental access.

Title: Seth Lloyd, Massachusetts Institute of Technology

Abstract to come

QTML Invited Talks

Title: AI for Quantum: Toward AI-Enhanced Quantum Computing Applications by Kohei Nakaji, NVIDIA

Abstract: The convergence of artificial intelligence (AI) and quantum computing represents one of the most promising frontiers in modern computational science. While ''quantum for AI'' has been widely explored as a potential application of quantum computing, ''AI for quantum'' — leveraging AI technologies to enhance quantum algorithms and quantum hardware — is now rapidly emerging, particularly with the advent of modern generative model techniques. In this talk, I will focus on how contemporary AI methods can be utilized to accelerate the development of quantum algorithms and enable next-generation quantum computing applications, including our proposed Generative Quantum Eigensolver (GQE) as a concrete example of such AI-driven approaches.

Title: Quantum Generative Modeling Beyond the NISQ Era by Michele Grossi, CERN

Abstract: Quantum computing provides a natural framework for generative modeling through sampling tasks with established complexity-theoretic advantages, yet standard parametrized-circuit approaches face persistent challenges in trainability and scalability. This talk reports recent progress on two complementary algorithmic directions developed to address these issues. The first is a differentiable quantum generative model (DQGM) based on quantum Chebyshev transforms, which enables post-training resolution scaling and efficient sampling without additional optimization. The second centers on quantum Boltzmann machines (QBMs), which offer a fault-tolerant path to scalable generative learning. A semi-quantum RBM (sqRBM) architecture with a commuting-visible Hamiltonian structure allows closed-form expressions for probabilities and gradients, providing provable expressive advantages over classical RBMs. Building on this, a quantum variant of contrastive divergence achieves O(1) forward-pass scaling for training. Theoretical results from these works are supported by numerical simulations, outlining a scalable and resource-efficient route for quantum generative modeling beyond the NISQ era.

Title: QuantumBoost: A lazy, yet fast, quantum algorithm for learning by Amira Abbas, Google

Abstract: The technique of combining multiple votes to enhance the quality of a decision is the core of boosting algorithms in machine learning. In particular, boosting provably increases decision quality by combining multiple "weak learners"—hypotheses that are only slightly better than random guessing—into a single "strong learner" that classifies data well. Inspired by work by Barak, Hardt and Kale, I will introduce QuantumBoost, a quantum algorithm that achieves the best known runtime complexity over other boosting methods. I will also share some interesting insights in the way my collaborators (Yanlin Chen, Tuyen Nguyen and Ronald de Wolf) and I ultimately developed QuantumBoost and proved its correctness with the help of Gemini's Deep Think model.

Title: The state of learning stabilizer-like states by Srinivasan Arunachalam, IBM

Abstract: This will be an overview talk wherein I go over the recent works in the last few years on learning stabilizer states (and their generalizations). I will discuss a couple of recent works wherein we give tolerant testing protocols for these states as well as applications to learning states with bounded stabilizer rank.

Title: Models of Learning for Quantum Processes: with noise, limitations and adversaries! by Mina Doosti, University of Edinburgh

Abstract: Characterizing a quantum system by learning its evolution is a fundamental problem with a myriad of applications. In this talk, I will explore different models for learning quantum processes from a quantum learning theory perspective, which are relevant in realistic scenarios involving noisy data or adversarial behaviour. These more-realistic quantum process learning models allow us to bridge the gap between sophisticated but contrived learning theory techniques and algorithms with provable guarantees, and real applications in areas such as physics and cryptography. I will discuss this importance and highlight two main physically motivated learning models: statistical query learning of quantum processes and agnostic process learning. Statistical queries are natural yet powerful learning models that provide both learning guarantees and robustness to noise [WD24,WD25]. Agnostic process learning [WLKD24], on the other hand, enables efficient learning of quantum processes even when the data source is noisy or potentially adversarial. Specifically, the model is formalized as follows: given query access to an unknown quantum channel Φ and a known concept class C of channels, the goal is to output a quantum

channel that approximates Φ as well as the best channel in C, up to some error. I will also discuss several natural applications of this model, including quantum machine learning, quantum metrology, classical simulation, and error mitigation. I will present relevant techniques and conclude with a discussion of open questions and limitations in both of these quantum process learning models.

Title: How to scale generative quantum machine learning to 1000 qubits by Joseph Bowles, Xanadu

Abstract: I will present a universal class of quantum generative models based on instantaneous quantum polynomial circuits. The training of these models can be performed entirely on classical hardware, allowing scaling to circuits with thousands of qubits and millions of parameters. By empirically implementing this training algorithm on large, real-world datasets, we will see how such models can be successfully trained despite the provable existence of barren plateaus under random parameter initialization. I will finish with a reflection regarding the outlook of the approach and the direction of the field in general.

Title: Shadows of quantum machine learning and shallow-depth learning separations by Sofiene Jerbi, Freie Universität Berlin

Abstract: In this talk, I will present two recent works related to the question of quantum advantages in machine learning. In the first work, we address a major obstacle to the widespread use of quantum machine learning models in practice: quantum models, even once trained, still require access to a quantum computer in order to be evaluated on new data. To solve this issue, we introduce a class of quantum models where quantum resources are only required during training, while the deployment of the trained model is classical. We prove that: (i) this class of models is universal for classically-deployed quantum machine learning; (ii) it does have restricted learning capacities compared to 'fully quantum' models, but nonetheless (iii) it achieves a provable learning advantage over fully classical learners; where (ii) and (iii) are contingent on widely believed assumptions in complexity theory. In the second work, we refine our understanding of the regimes where quantum advantages arise in machine learning, by proving a PAC learning advantage in the realm of shallow-depth circuits. This learning advantage has the particularity that it is unconditional, meaning that we do not need to make assumptions such as the existence of classically hard, quantumly easy, cryptographic functions to show an advantage.

QTML Highlighted Talks

Title: Quantum computing and persistence in TDA

Abstract: Topological data analysis (TDA) aims to extract noise-robust features of a data set by studying the number and persistence of holes in its topology. We show that a central task of TDA -- deciding whether a given hole persists across different length scales – is BQP_1-hard and contained in BQP, implying an exponential quantum speedup for this task under standard complexity-theoretic assumptions. Our results are based on the observation that the persistence of a hole can be encoded in the guided sparse Hamiltonian problem, where the guiding state is constructed from a harmonic representative of the hole.

Authors: Casper Gyurik, Alexander Schmidhuber, Robbie King, Vedran Dunjko and Ryu Hayakawa

Title: Do you know what q-means?

Abstract: "Clustering is one of the most important tools for analysis of large datasets, and perhaps the most popular clustering algorithm is Lloyd's algorithm for k-means. This algorithm takes n vectors in a d-dimensional space and outputs k centroid vectors, which partition the vectors into clusters based on which centroid is closest to a particular vector. We present a classical epsilon-k-means algorithm that performs an approximate version of one iteration of Lloyd's algorithm, with a time complexity that improves exponentially in the data size n compared to previous classical algorithms. It matches the runtime of the q-means quantum algorithm originally proposed by Kerenidis, Landman, Luongo, and Prakash (NeurIPS 2019). To our knowledge, this is the fastest classical algorithm for approximate k-means. We then turn our attention to the quantum setting, and propose an improved version of the q-means algorithm. Our new quantum algorithm achieves a better runtime than previous quantum approaches and offers a polynomial improvement over our classical epsilon-k-means in several parameters. Unlike prior quantum algorithms, our method does not rely on quantum linear algebra primitives. Instead, it uses QRAM to prepare simple quantum states based on the current cluster assignments and applies multivariate quantum amplitude estimation. Finally, we provide the first quantum and classical lower bounds for performing a single iteration of the k-means problem. These results show that our algorithms are optimal in most of the relevant parameters."

Authors: Arjan Cornelissen, Joao F. Doriguello, Alessandro Luongo and Ewin Tang

Title: A Bit of Freedom Goes a Long Way: Quantum and Classical Algorithms for Online Learning of MDPs under a Generative Model

Abstract: We propose novel classical and quantum online algorithms for learning finite-horizon and infinite-horizon average-reward Markov Decision Processes (MDPs). Our algorithms are based on a hybrid exploration-generative reinforcement learning (RL) model wherein the agent can, from time to time, freely interact with the environment in a generative sampling fashion, i.e., by having access to a ``simulator''. By employing known classical and new quantum algorithms for approximating optimal policies under a generative model within our learning algorithms, we show that it is possible to avoid several paradigms from RL like ``optimism in the face of uncertainty'' and ``posterior sampling'' and instead compute and use optimal policies directly, which yields better regret bounds compared to previous works. For finite-horizon MDPs, our quantum algorithm obtains regret bounds which only depend logarithmically on the number of time steps T, thus breaking the $O(\sqrt{T})$ classical barrier. This matches the time dependence of the prior quantum works of Ganguly et al.~(arXiv'23) and Zhong et al.~(ICML'24), but with improved dependence on other parameters like state space size S and action space size A. For infinite-horizon MDPs, our classical and quantum bounds still maintain the $\widetilde{O}(\sqrt{T})$ dependence but with better S and A factors. Nonetheless, we propose a novel measure of regret for infinite-horizon MDPs with respect to which our quantum algorithm has $\poly\log{T}$ regret, exponentially better compared to classical algorithms. Finally, we generalise all of our results to compact continuous state spaces.

Authors: Andris Ambainis, Joao F. Doriguello and Debbie Huey Chih Lim

Title: Decoded Quantum Interferometry

Abstract: Whether quantum computers can achieve exponential speedups in optimization has been a major open question in quantum algorithms since the field began. Here we introduce a quantum algorithm called Decoded Quantum Interferometry (DQI), which uses the quantum

Fourier transform to reduce optimization problems to decoding problems. For approximating optimal polynomial fits to data over finite fields, DQI efficiently achieves a better approximation ratio than any polynomial time classical algorithm known to us, thus suggesting exponential quantum speedup. Sparse unstructured optimization problems such as max-k-XORSAT are reduced to decoding of LDPC codes. We prove a theorem which allows the performance of DQI to be calculated instance-by-instance based on the empirical performance of classical decoders.
We use this to construct an instance of max-XORSAT for which DQI finds an approximate
optimum that cannot be found by simulated annealing or any of the other general-purpose classical heuristics that we tried, unless given five orders of magnitude more compute time than the decoding problem requires. Although we subsequently design a tailored classical solver that beats DQI within reasonable runtime, our results nevertheless demonstrate that the combination of quantum Fourier transforms with powerful decoding primitives provides a promising new approach to finding quantum speedups for hard optimization problems.
This submission presents the original DQI algorithm alongside recent improvements and generalizations.
Authors: Stephen Jordan, Noah Shutty, Mary Wootters, Adam Zalcman, Alexander Schmidhuber, Robbie King, Sergei Isakov, Tanuj Khattar and Ryan Babbush

QTML Accepted Talks

Title: Quartic quantum speedups for planted inference
Abstract: We describe a quantum algorithm for the Planted~Noisy~kXOR problem (also known as sparse Learning Parity with Noise) that achieves a nearly quartic (4th power) speedup over the best known classical algorithm while also only using exponentially less space. Our work generalizes and simplifies prior work of Hastings, by building on his quantum algorithm for the Tensor Principal Component Analysis (PCA) problem. We achieve our quantum speedup using a general framework based on the Kikuchi Method (recovering the quartic speedup for Tensor PCA), and we anticipate it will yield similar speedups for other planted inference problems. These speedups rely on the fact that planted inference problems naturally instantiate the Guided Sparse Hamiltonian problem. Since planted inference problems serve as a testbed for studying the hardness of statistical learning, our work paves a path towards significant polynomial quantum speedups in machine learning.
Authors: Alexander Schmidhuber, Ryan O'Donnell, Robin Kothari and Ryan Babbush

Title: Learning pure quantum states (almost) without regret
Abstract: We initiate the study of quantum state tomography with minimal disturbance to the samples. Can we learn a precise description of a quantum state through sequential measurements of samples while at the same time making sure that the post-measurement state of the samples is only minimally perturbed? Defining regret as the cumulative disturbance of all samples, the challenge is to find a balance between the most informative sequence of measurements on the one hand and measurements incurring minimal regret on the other. Here we answer this question for pure qubit states by exhibiting a protocol that achieves maximal precision while incurring a regret that grows only polylogarithmically with the number of samples, a scaling that we show to be optimal.
Authors: Josep Lumbreras, Mikhail Terekhov and Marco Tomamichel

Title: Nearly query-optimal classical shadow estimation of unitary channels

Abstract: Classical shadow estimation (CSE) is a powerful tool for learning the properties of quantum states and quantum processes. Here we consider the CSE task for quantum unitary channels. Based on collective measurements on multiple systems, we propose a query efficient protocol for this task, whose query complexity has a quadratic advantage over the previous best approach for this problem, and almost saturates the information-theoretic lower bound. To further enhance practicality, we present a variant protocol using only single-copy measurements, which still offers much better query performance than previous protocols that do not use quantum memories. This protocol can also serve as a key subroutine for learning an arbitrary unknown Hamiltonian from dynamics, outperforming existing approaches to this problem.

Authors: Zihao Li, Changhao Yi, You Zhou and Huangjun Zhu

Title: Multiple-basis representation of quantum states

Abstract: Classical simulation of quantum physics is a central approach to investigating physical phenomena. Quantum computers enhance computational capabilities beyond those of classical resources, but it remains unclear to what extent existing limited quantum computers can contribute to this enhancement.

In this work, we explore a new hybrid, efficient quantum-classical representation of quantum states, the multiple-basis representation. This representation consists of a linear combination of states that are sparse in some given bases, specified by quantum circuits.

Such representation is particularly appealing when considering depth-limited quantum circuits within reach of current hardware.

We analyze the expressivity of multiple-basis representation states depending on the classical simulability of their quantum circuits. In particular, we show that multiple-basis representation states include, but are not restricted to, both matrix-product states and stabilizer states.

Furthermore, we investigate applications of this representation in the problems of approximation of ground states, simulation of deeper computations by specifying bases with shallow circuits, and a tomographical protocol to describe states as multiple-basis representations.

We envision this work to open the path of simultaneous use of several hardware-friendly bases, a natural description of hybrid computational methods accessible for near-term hardware.

Authors: Adrián Pérez-Salinas, Patrick Emonts, Jordi Tura Brugués and Vedran Dunjko

Title: Verifiable End-to-End Delegated Variational Quantum Algorithms

Abstract: Variational quantum algorithms (VQAs) have emerged as promising candidates for solving complex optimization and machine learning tasks on near-term quantum hardware.

However, due to hardware limitations, small-scale users face challenges executing quantum operations, making delegation to more powerful quantum devices desirable. In this work, we introduce a framework for delegated variational quantum algorithms (DVQAs), where a client with limited quantum capabilities delegates the execution of a VQA to a more powerful quantum server. In particular, we introduce a protocol that enables a client to delegate a variational quantum algorithm to a server while ensuring that the input, the output and also the computation itself remain secret. Additionally, if the protocol does not abort, the client can be certain that the computation outcome is indeed correct. Our approach first proposes a verifiable Protocol for delegating the quantum computation required at each optimization step of a VQA, and then combines the iterative steps into an error-resilient optimization process that offers end-to-end verifiable algorithm execution. Our results demonstrate that secure delegation of variational

quantum algorithms is a realistic solution for near-term quantum networks, paving the way for practical quantum cloud computing applications.
Authors: Matteo Antonio Inajetovic, Petros Wallden and Anna Pappa

Title: Efficient quantum-enhanced classical simulation for patches of quantum landscapes

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

Title: On the dynamical Lie algebras of quantum approximate optimization algorithms

Abstract: Dynamical Lie algebras (DLAs) have emerged as a valuable tool in the study of parameterized quantum circuits, helping to characterize both their expressiveness and
trainability. In particular, the absence or presence of barren plateaus (BPs)---flat regions in parameter space that prevent the efficient training of variational quantum algorithms---has recently been shown to be intimately related to quantities derived from the associated DLA.
In this work, we investigate DLAs for the quantum approximate optimization algorithm (QAOA), one of the most studied variational quantum algorithms for solving graph MaxCut and other combinatorial optimization problems. While DLAs for QAOA circuits have been studied before, existing results have either been based on numerical evidence, or else correspond to DLA
generators specifically chosen to be universal for quantum computation on a subspace of states. We initiate an analytical study of barren plateaus and other statistics of QAOA algorithms, and give bounds on the dimensions of the corresponding DLAs and their centers for general graphs. We then focus on the n-vertex cycle and complete graphs. For the cycle graph we give an explicit basis, identify its decomposition into the direct sum of a 2-dimensional center and a semisimple component isomorphic to n-1 copies of su(2). We give an explicit basis for this
isomorphism, and a closed-form expression for the variance of the cost function, proving the absence of BPs. For the complete graph we prove that the dimension of the DLA is $O(n^3)$ and give an explicit basis for the DLA.
Authors: Jonathan Allcock, Miklos Santha, Pei Yuan and Shengyu Zhang

Title: Shedding light on classical shadows

Abstract: In this work, we introduce a shadow tomography protocol tailored to linear optical systems. Our protocol enables the tomography of number states—and superpositions thereof—using only passive linear optical transformations and photon-number resolving (PNR) detectors. Extra technicalities emerge as this setting is not tomographic complete but is practically

relevant due to its experimental accessibility. We adapt the classical shadow framework to this context and provide both sample and time complexity bounds.
In particular, we characterize the class of observables that can be estimated efficiently, leveraging the visible space formalism. Our analysis reveals that, even under the constraints of passive linear optics, useful and scalable shadow tomography is possible, opening the door to new practical applications in photonic quantum computing.
Authors: Hugo Thomas, Pierre-Emmanuel Emeriau and Ulysse Chabaud

Title: Quantum simulation with sum-of-squares spectral amplification

Abstract: We present sum-of-squares spectral amplification (SOSSA), a framework for improving quantum simulation relevant to low-energy problems. We show how SOSSA can be applied to energy and phase estimation and provide fast quantum algorithms for these problems that significantly improve over prior art. To illustrate the power of SOSSA in applications, we consider the Sachdev-Ye-Kitaev model, a representative strongly correlated system, and demonstrate asymptotic speedups over generic simulation methods by a factor that is the square root of the system size. Our results reinforce those observed in [G.H. Low \textit{et al.}, arXiv:2502.15882 (2025)], where SOSSA was used to achieve state-of-the-art gate costs for phase estimation of real-world quantum chemistry systems.
Authors: Nicholas Rubin, Guanghao Low, Robbie King, Eugene DePrince, Alec White, Ryan Babbush, Dominic Berry and Rolando Somma

Title: Variational quantum algorithms with exact geodesic transport

Abstract: Variational quantum algorithms (VQAs) are promising candidates for near-term applications of quantum computers, but their training represents a major challenge in practice. We introduce exact-geodesic VQAs, a space-curvature aware framework that enables analytic Riemannian optimization of variational quantum circuits through a convenient choice of circuit ansatz. Our method exploits the exact metric to find a near-optimal parameter optimization path based on exact geodesic transport with conjugate gradients (EGT-CG). This supersedes the celebrated quantum natural gradient method, in fact recovering it as its first-order approximation. Further, the exact-geodesic updates for our circuit ansatz have the same cost as standard gradient descent. This contrasts with previous metric-aware methods, which require resource-intensive estimations of the metric. For chemistry problems of up to 14 electrons, our toolkit allows us to achieve up to a 20x reduction in the number of iterations over Adam or quantum natural gradient methods. Moreover, for degenerate problems, which are notoriously difficult to optimize with conventional methods, we achieve rapid convergence to the global minima. Our work demonstrates that the cost of VQA optimization can be drastically reduced by harnessing the Riemannian geometry of the manifold expressed by the circuit ansatz, with both practical and fundamental implications at the interface between quantum machine learning, differential geometry, and optimal control theory.
Authors: André Ferreira-Martins, Renato M. S. Farias, Giancarlo Camilo, Thiago O. Maciel, Allan Tosta, Ruge Lin, Abdulla Alhajri, Tobias Haug and Leandro Aolita

Title: A unifying account of warm start guarantees for patches of quantum landscapes

Abstract: Barren plateaus are fundamentally a statement about quantum loss landscapes on average but there can, and generally will, exist patches of barren plateau landscapes with substantial gradients. Previous work has studied certain classes of parameterized quantum circuits

and found example regions where gradients vanish at worst polynomially in system size. Here we present a general bound that unifies all these previous cases and that can tackle physically-motivated ansätze that could not be analyzed previously. Concretely, we analytically prove a lower-bound on the variance of the loss that can be used to show that in a non-exponentially narrow region around a point with curvature the loss variance cannot decay exponentially fast. This result is complemented by numerics and an upper-bound that suggest that any loss function with a barren plateau will have exponentially vanishing gradients in any constant radius subregion. Our work thus suggests that while there are hopes to be able to warm-start variational quantum algorithms, any initialization strategy that cannot get increasingly close to the region of attraction with increasing problem size is likely inadequate.
Authors: Hela Mhiri, Ricard Puig, Manuel Rudolph, Sacha Lerch, Thiparat Chotibut, Supanut Thanasilp and Zoe Holmes

Title: Hamiltonian Locality Testing via Trotterized Postselection

Abstract: The (tolerant) Hamiltonian locality testing problem, Introduced in [Bluhm, Caro,Oufkir `24], is to determine whether a Hamiltonian H is ε_1-close to being k-local (i.e., can be written as the sum of weight-k Pauli operators) or ε_2-far from any k-local Hamiltonian, given access to its time evolution operator and using as little total evolution time as possible, with distance typically defined by the normalized Frobenius norm. We give the tightest known bounds for this problem, proving an $O(\sqrt{\frac{\eps_2}{(\eps_2-\eps_1)^5}})$ evolution time upper bound and an $\Omega(\frac{1}{\eps_2-\eps_1})$ lower bound. Our algorithm does not require reverse time evolution or controlled application of the time evolution operator, although our lower bound applies to algorithms using either tool. Furthermore, we show that if we are allowed reverse time evolution, this lower bound is tight, giving a matching $O(\frac{1}{\eps_2-\eps_1})$ evolution time algorithm.
Authors: John Kallaugher and Daniel Liang

Title: A Unified Theory of Quantum Neural Network Loss Landscapes

Abstract: Classical neural networks with random initialization famously behave as Gaussian processes in the limit of many neurons, which allows one to completely characterize their training and generalization behavior. No such general understanding exists for quantum neural networks (QNNs), which---outside of certain special cases---are known to not behave as Gaussian processes when randomly initialized. We here prove that QNNs and their first two derivatives instead generally form what we call Wishart processes, where certain algebraic properties of the network determine the hyperparameters of the process. This Wishart process description allows us to, for the first time: give necessary and sufficient conditions for a QNN architecture to have a Gaussian process limit; calculate the full gradient distribution, generalizing previously known barren plateau results; and calculate the local minima distribution of algebraically constrained QNNs. Our unified framework suggests a certain simple operational definition for the "trainability" of a given QNN model using a newly introduced, experimentally accessible quantity we call the degrees of freedom of the network architecture.
Authors: Eric Anschuetz

Title: Quantum HodgeRank: Topology-Based Rank Aggregation on Quantum Computers

Abstract: HodgeRank generalizes ranking algorithms, e.g. Google PageRank, to rank alternatives based on real-world (often incomplete) data using graphs and discrete exterior calculus. It analyzes multipartite interactions on high-dimensional networks with a complexity that scales exponentially with dimension. We develop a quantum algorithm that approximates the HodgeRank solution with complexity independent of dimension. Our algorithm extracts relevant information from the state such as the ranking consistency, which achieves a superpolynomial speedup over similar classical methods.

Authors: Caesnan Leditto, Angus Southwell, Behnam Tonekaboni, Muhammad Usman and Kavan Modi

Title: Interactive proofs for verifying (quantum) learning and testing

Abstract: We consider the problem of testing and learning from data in the presence of resource constraints, such as limited memory or weak data access, which place limitations on the efficiency and feasibility of testing or learning. In particular, we ask the following question: Could a resource-constrained learner/tester use interaction with a resource-unconstrained but untrusted party to solve a learning or testing problem more efficiently than they could without such an interaction? In this work, we answer this question both abstractly and for concrete problems, in two complementary ways: For a wide variety of scenarios, we prove that a resource-constrained learner cannot gain any advantage through classical interaction with an untrusted prover. As a special case, we show that for the vast majority of testing and learning problems in which quantum memory is a meaningful resource, a memory-constrained quantum algorithm cannot overcome its limitations via classical communication with a memory-unconstrained quantum prover. In contrast, when quantum communication is allowed, we construct a variety of interactive proof protocols, for specific learning and testing problems, which allow memory-constrained quantum verifiers to gain significant advantages through delegation to untrusted provers. These results highlight both the limitations and potential of delegating learning and testing problems to resource-rich but untrusted third parties.

Authors: Matthias C. Caro, Jens Eisert, Marcel Hinsche, Marios Ioannou, Alexander Nietner and Ryan Sweke

Title: Learning Quantum States with Tunable Loss Functions

Abstract: Learning from quantum data presents new challenges to the paradigm of learning from data. This typically entails the use of quantum learning models to learn quantum processes, that come with enough subtleties to modify the theoretical learning frameworks. This intersection warrants new frameworks for complexity measures, including those on quantum sample complexity and generalization bounds. Empirical risk minimization (ERM) serves as the foundational framework for evaluating learning models in general. The necessity for regularization strategies leads to the development of advanced regularization strategies such as tilted empirical risk minimization (TERM). Theoretical aspects of quantum learning under a quantum ERM framework are presented in [PRX Quantum 5, 020367 (2024)]. In this work, we propose a definition for TERM suitable to be employed when learning quantum processes, which gives rise to quantum TERM (QTERM). By extension, QTERM can be viewed as a regularization strategy for quantum state learning. This work contributes to the existing literature on quantum and classical physics threefold. First, we prove QTERM learnability by deriving upper bounds on QTERM's sample complexity. Second, we establish new PAC generalization bounds on classical TERM. Third, we present QTERM agnostic learning guarantees for quantum hypothesis selection. These results

contribute to the broader literature of complexity bounds on the feasibility of learning quantum states, as well as the presence of regularization techniques in quantum learning.
Authors: Yixian Qiu, Lirandë Pira and Patrick Rebentrost

Title: An efficient approach to realize Quantum Random Features

Abstract: It has recently been shown that several quantum machine learning models, including Quantum Circuit Learning (QCL), Quantum Reservoir Computing (QRC), and Quantum Extreme Learning Machines (QELM), can be analyzed via the Fourier series. In such works, the data-encoding circuit determines the accessible frequency components, and increasing circuit depth expands the representational capacity. However, these models face challenges in practical applications due to the lack of control over which frequency components are needed for a given task. In our talk, we propose a practical QRC/QELM architecture inspired by Random Fourier Features (RFF). Our model utilizes layered quantum circuits with Z-rotation encoders and fixed permutation unitaries to efficiently generate RFF-like frequency structures. We demonstrate that our method requires only $\mathcal{O}(log N_f\cdot L)$ computational cost to generate N_f features in preprocessing, compared to $\mathcal{O}(N_f)$ in classical RFF. Furthermore, the model remains robust when the permutation circuit is replaced with more general quantum dynamics. These results provide practical design principles for constructing expressive and scalable QML models, applicable to tasks such as image classification.
Authors: Akitada Sakurai, Aoi Hayashi, William Munro and Kae Nemoto

Title: Kernel-based Dequantization of Quantum Machine Learning

Abstract: A key challenge in quantum machine learning (QML) is to identify tasks where quantum models offer a genuine advantage over classical approaches. One promising strategy is dequantization—showing that classical algorithms can match the performance of quantum models for specific tasks. In this work, we study classical kernel-based methods as dequantization methods for QML models based on parameterized quantum circuits (PQCs), including quantum neural networks and quantum kernel methods. For a given PQC, we define a corresponding family of trigonometric kernels that can capture the function class expressible by the quantum model. We approximate these kernels using Random Fourier Features (RFF) and derive theoretical bounds on the difference between the true risk of RFF-approximated classical models and their quantum counterparts in both regression and classification tasks. Based on these bounds, we identify sufficient conditions for dequantization. Finally, we show that in certain cases, the exact kernel can be computed efficiently via tensor networks, removing the need for approximation. Our results can be used in the design of PQC-based QML models to avoid dequantization. Moreover, our proposed dequantization schemes can be used as independent heuristic classical algorithms before deploying costly QML models.
Authors: Alice Barthe, Jens Eisert, Bryce Fuller, Elies Gil-Fuster, Michele Grossi, Zoë Holmes, Sofiène Jerbi, Johannes Jakob Meyer, Erik Recio-Armengol, Mehrad Sahebi, Seongwook Shin, Yudai Suzuki and Ryan Sweke

Title: A good basis allows for classical shadows with arbitrary group representations

Abstract: Classical Shadows (CSs) have emerged as an efficient method to predict many properties of an unknown quantum state. A common denominator in CS protocols is to parametrize the random measurement bases with a random group action. While many different groups have been proposed, they have been mostly studied on a case-by-case basis, like random Paulis and Cliffords, or matchgates. In an step toward a unified analysis, recent work showed that one can obtain a general expression for the shadows channel corresponding to the action of an arbitrary group G, provided each irreducible representation appears at most once. The importance of this realization is that it allows one to access a whole family of CS protocols parametrized by the choice of group representations. Moreover, the channel can be inverted analytically, i.e., at no cost. However, many—if not most—relevant settings involve multiplicities and thus remain beyond the scope of previous analyses.

In this work, we show one can go beyond this multiplicity limitation and obtain an analytically invertible channel on arbitrary group representations. We give an explicit choice of basis which we call a ``good basis'' such that the measurement channel becomes a weighted sum of projectors into the irreducible representation. We complement this analysis by deriving general bounds on the variance of the estimators which directly relate to sample-complexity bounds, and discuss how to implement good-basis measurements.

Our method both unifies existing CS procedures based on local, global, orthogonal, symplectic and fermionic Gaussian unitaries, and allows us to easily generate new protocols based on other groups, or different representations of previous groups. For example, we characterize novel shadow protocols based on sampling from the orthogonal group, the spin and tensor representations of SU(2), and the exceptional Lie group G2.

Authors: Martin Larocca, Maxwell West, Marco Cerezo, Frederic Sauvage, Roy Forestano, Nathan Killoran and David Wierichs

Title: Bayesian Quantum Orthogonal Neural Networks for Anomaly Detection

Abstract: Identification of defects or anomalies in 3D objects is a crucial task to ensure correct functionality. In this work, we combine Bayesian learning with recent developments in quantum and quantum-inspired machine learning, specifically orthogonal neural networks, to tackle the anomaly detection problem for an industrially relevant use case. Bayesian learning enables uncertainty quantification of predictions, while orthogonality in weight matrices enables smooth training. We develop orthogonal (quantum) versions of 3D convolutional neural networks and show that these models can successfully detect anomalies in 3D objects. To test the feasibility of incorporating quantum computers into a quantum-enhanced anomaly detection pipeline, we perform hardware experiments with our models on IBM's 127 qubit Brisbane device, testing the effect of noise and limited measurement shots.

Authors: Natansh Mathur, Brian Coyle, Nishant Jain, Snehal Raj, Akshat Tandon, Jasper Krauser and Rainer Stoessel

Title: Polynomial Speed-Up in Photonic Neural Networks via Adaptive State Injection

Abstract: Linear optical architectures have been extensively investigated for quantum computing and quantum machine learning applications. Recently, proposals for photonic quantum machine learning have combined linear optics with resource adaptivity, such as adaptive circuit reconfiguration, which promises to enhance expressivity and improve algorithm performances and scalability. Moreover, linear optical platforms preserve some subspaces due to the fixed number of particles during the computation, a property recently exploited to design a novel quantum convolutional neural networks. This last architecture has shown an advantage in terms of running time complexity and of the number of parameters needed with respect to other quantum neural network proposals. We propose to present the results from two papers of our team. First, we propose a new scheme that relies on state injection, a measurement-based technique that can produce states that are more controllable, and solve learning tasks that are believed to be intractable classically. Secondly, we design and experimentally implement the first photonic

quantum convolutional neural network (PQCNN) architecture based on particle-number preserving circuits equipped with state injection. Subsequently, we experimentally validate the PQCNN for an image classification on a photonic platform utilizing a semiconductor quantum dot-based single-photon source and programmable integrated photonic interferometers comprising 8 and 12 modes. We highlight the potential utility of a simple adaptive technique for a nonlinear Boson Sampling task, compatible with near-term quantum devices. Such approach open the path to new QML algorithms with useful polynomial advantages, as such photonic architecture present a very low running time.
Authors: Léo Monbroussou, Beatrice Polacchi, Verena Yacoub, Eliott Z. Mamon, Hugo Thomas, Eugenio Caruccio, Giovanni Rodari, Francesco Hoch, Gonzalo Carvacho, Nicolo Spagnolo, Taira Giordani, Mattia Bossi, Abhiram Rajan, Niki Di Giano, Riccardo Albiero, Francesco Ceccarelli, Roberto Osellame, Ulysse Chabaud, Fabio Sciarrino and Elham Kashefi

Title: Natural gradient and parameter estimation for quantum Boltzmann machines
Abstract: Thermal states play a fundamental role in various areas of physics, and they are becoming increasingly important in quantum information science, with applications related to semi-definite programming, quantum Boltzmann machine learning, Hamiltonian learning, and the related task of estimating the parameters of a Hamiltonian. Here we establish formulas underlying the basic geometry of parameterized thermal states, and we delineate quantum algorithms for estimating the values of these formulas. More specifically, we prove formulas for the Fisher–Bures and Kubo–Mori information matrices of parameterized thermal states, and our quantum algorithms for estimating their matrix elements involve a combination of classical sampling, Hamiltonian simulation, and the Hadamard test. These results have applications in developing a natural gradient descent algorithm for quantum Boltzmann machine learning, which takes into account the geometry of thermal states, and in establishing fundamental limitations on the ability to estimate the parameters of a Hamiltonian, when given access to thermal-state samples. For the latter task, and for the special case of estimating a single parameter, we sketch an algorithm that realizes a measurement that is asymptotically optimal for the estimation task. We finally stress that the natural gradient descent algorithm developed here can be used for any machine learning problem that employs the quantum Boltzmann machine ansatz.
Authors: Dhrumil Patel and Mark Wilde

Title: A PAC-Bayesian approach to generalization for quantum models
Abstract: Generalization is a central concept in machine learning, yet for quantum models, it is predominantly analyzed through uniform bounds that depend on a model's overall capacity rather than the specific function learned. These bounds are often too loose and insensitive to the training process. In this work, we address this limitation by deriving the first PAC-Bayesian generalization bound for a broad class of quantum machine learning (QML) models. Our framework analyzes layered circuits composed of general quantum channels, parameterized via their process matrices, which include unitary, dissipative, and feedforward operations. By performing a channel perturbation analysis, we establish a non-uniform generalization bound that explicitly depends on the norms of the post-training parameter matrices. This data-dependent nature allows our bound to reflect the properties of the learned solution, and we show that it can offer improvements over existing uniform, covering-number-based bounds. By connecting channel perturbation theory with the powerful PAC-Bayesian framework, our work provides a foundational tool for a more nuanced, training-aware analysis of generalization in QML.

Authors: Pablo Rodriguez-Grasa, Matthias C. Caro, Jens Eisert, Elies Gil-Fuster, Franz J. Schreiber and Carlos Bravo-Prieto

Title: Fourier Fingerprints of Ansatzes in Quantum Machine Learning

Abstract: The training of parameterized quantum circuits (PQCs) as machine learning models is one of the most widely studied paradigms in quantum machine learning (QML). In this framework, a typical model consists of quantum feature maps, which encode a classical input into the Hilbert space, and variational ansatzes which manipulate the input via trainable parameterized gates.
Importantly, the output of these models can be represented as a partial Fourier series in the input, giving rise to the name quantum Fourier models (QFMs). Fully understanding how the choice of feature map and ansatz affects the Fourier spectrum of QFMs is crucial for maximizing their performance, including for popular frameworks such as quantum neural networks (QNNs) and quantum kernels. In this work, we theoretically motivate the appearance of correlations between the Fourier coefficients of QFMs, implying the inability to control each term in the Fourier series independently. For a range of popular ansatzes, we compute these correlations and find a unique pattern for each ansatz, which we call the Fourier fingerprint. Subsequently, we demonstrate the utility of the Fourier coefficient correlation (FCC), derived from the fingerprint, for learning random 1D and 2D Fourier series. In these experiments, we show that the FCC can effectively predict which ansatzes will yield the best performance. Finally, we demonstrate how our the FCC applies to the more challenging problem of 2D jet reconstruction in high-energy physics and find that the FCC can again be used to predict the best performing ansatz before training occurs. Overall, our results illustrate that the Fourier fingerprint is a powerful new tool for the problem of optimal ansatz choice in QML.
Authors: Melvin Strobl, Emre Sahin, Lucas van der Horst, Eileen Kuehn, Achim Streit and Ben Jaderberg

Title: Efficient learning for linear properties of bounded-gate quantum circuits

Abstract: The vast and complicated many-qubit state space forbids us to comprehensively capture the dynamics of modern quantum computers via classical simulations or quantum tomography. Recent progress in quantum learning theory prompts a crucial question: can linear properties of a many-qubit circuit with d tunable RZ gates and G-d Clifford gates be efficiently learned from measurement data generated by varying classical inputs? In this work, we prove that the sample complexity scaling linearly in d is required to achieve a small prediction error, while the corresponding computational complexity may scale exponentially in d. To address this challenge, we propose a kernel-based method leveraging classical shadows and truncated trigonometric expansions, enabling a controllable trade-off between prediction accuracy and computational overhead. Our results advance two crucial realms in quantum computation: the exploration of quantum algorithms with practical utilities and learning-based quantum system certification. We conduct numerical simulations to validate our proposals across diverse scenarios, encompassing quantum information processing protocols, Hamiltonian simulation, and variational quantum algorithms up to 60 qubits.
Authors: Yuxuan Du, Min-Hsiu Hsieh and Dacheng Tao

Title: Scalable Neural Decoders for Practical Real-Time Quantum Error Correction

Abstract: Implementing efficient and scalable decoders for quantum error correction is essential for practical quantum computing. Recurrent transformer-based architectures such as AlphaQubit

achieve high decoding accuracy but suffer from prohibitive computational costs. To address this, we introduce a Mamba decoder that replaces each Multi-Head Attention block of AlphaQubit with a Mamba module. On Google's Sycamore memory experiment, our Mamba decoder matches transformer-level performance, achieving logical error rates of 2.98×10^{-2} at distance 3 and 3.03×10^{-2} at distance 5. We further evaluate real-time performance over 400 cycles with a latency-dependent noise model tied to computational complexity. The transformer's prohibitive $O(d^4)$ complexity leads to a severe accumulation of decoder-induced errors, whereas the Mamba decoder's efficient $O(d^2)$ scaling avoids this problem, demonstrating more robust performance. Our results thus highlight Mamba's superior speed-accuracy trade-off, establishing it as a viable architecture for large-scale, real-time decoders for quantum error correction.

Authors: Changwon Lee, Tak Hur and Daniel Kyungdeock Park

Title: Pauli Propagation: A computational framework for simulating quantum systems

Abstract: Pauli propagation (PP) is a relative newcomer to the corpus of classical simulation algorithms and yet is already competitive with other state-of-the-art methods for certain tasks. At their core, PP methods approximate the evolution of a quantum operator (typically, an observable in the Heisenberg picture) via a truncated Pauli path integral. This approach tends to be limited by very different quantum circuit characteristics compared to, for example, popular and potent tensor network methods.

We provide a comprehensive account of this new simulation framework and present PauliPropagation.jl: the first general-purpose open-source Pauli propagation software. We discuss the Pauli propagation framework from high to low level, from the algorithmic formulation and known theoretical results, to our lived experience and practical implementation details. Furthermore, we present two theoretical efficiency guarantees for PP simulation of noise-free scrambling circuits and circuits subject to arbitrary local noise. We hope this new tool not only proves useful for studying quantum systems and benchmarking quantum hardware, but also inspires new classical-quantum frameworks utilizing the best and most suitable classical simulation method in conjunction with upcoming quantum computers.

Authors: Manuel Rudolph, Armando Angrisani, Tyson Jones, Yanting Teng, Alexander Schmidhuber, Antonio Anna Mele, Marco Cerezo, Hsin-Yuang Huang and Zoe Holmes

Title: Designing quantum machine learning models for graphs

Abstract: Geometric Machine Learning (GML) successes have been achieved through thorough study and design of new equivariant neural networks. In comparison, geometric quantum machine learning (GQML) models critically lacks such a detailed understanding and a unifying perspective on their design remains elusive. Focusing on GQML models for graph datasets, we show that a comprehensive characterisation of their constituents is possible. We further probe benefits of this toolbox including the generalization of known models, sometimes at virtually no cost, and straightforward classical pre-training strategies.

Authors: Frederic Sauvage, Pranav Kalidindi, Frederic Rapp and Martin Larocca

Title: Sample importance in data-driven decoding

Abstract: Data-driven decoding (DDD) -- learning to decode syndromes of (quantum) error-correcting codes using training examples -- can be a difficult problem due to several atypical and poorly understood properties of the training data. We introduce a theory of example

importance that clarifies these unusual aspects of DDD: For instance, we show that DDD of a simple error-correcting code is equivalent to a noisy, imbalanced binary classification problem. This suggests that an existing data augmentation technique – turning the knob to increase error rates in training data -- actually introduces a tradeoff between class imbalance and label noise. We apply this technique in experiments showing robust improvements to decoder accuracy while
explaining the failures of this technique in terms of example importance. We show similar improvements for decoding quantum codes involving multiple rounds of syndrome measurements and we characterize example importance in random stabilizer codes, suggesting broad applicability of both example importance and turning the knob for improving experimentally relevant data-driven decoders.
Authors: Evan Peters

Title: Productionizing Quantum Mass Production

Abstract: For many practical applications of quantum computing, the most costly steps involve coherently accessing classical data. We help address this challenge by applying mass
production techniques, which can reduce the cost of applying an operation multiple times in parallel. We combine these techniques with modern approaches for classical data loading based on "quantum read-only memory" (QROM). We find that we can polynomially reduce the total
number of gates required for data loading, but we find no advantage in cost models that only count the number of non-Clifford gates. Furthermore, for realistic cost models and problem sizes, we find that it is possible to reduce the cost of parallel data loading by an order of magnitude
or more. We present several applications of quantum mass production, including a scheme that uses parallel phase estimation to asymptotically reduce the gate complexity of state-of-the-art algorithms for estimating eigenvalues of the quantum chemical Hamiltonian. We also show that mass production can be used to reduce the cost of serial calls to the same data loading oracle by precomputing several copies of a novel QROM resource state.
Authors: William J. Huggins, Tanuj Khattar and Nathan Wiebe

Title: Mildly-Interacting Fermionic Unitaries are Efficiently Learnable

Abstract: Recent work has shown that one can efficiently learn fermionic Gaussian unitaries, also commonly known as nearest-neighbor matchcircuits or non-interacting fermionic unitaries. However, one could ask a similar question about unitaries that are near Gaussian: for example, unitaries prepared with a small number of non-Gaussian circuit elements. These operators find significance in quantum chemistry and many-body physics, yet no algorithm exists to learn them.
We give the first such result by devising an algorithm which makes queries to an n-mode fermionic unitary U prepared by at most $O(t)$ non-Gaussian gates and returns a circuit approximating U to diamond distance ε in time $poly(n,2^t,1/\varepsilon)$. This
resolves a central open question of Mele and Herasymenko under the strongest distance metric. In fact, our algorithm is much more general: we define a property of unitary Gaussianity known as unitary Gaussian dimension and show that our algorithm can learn n-mode unitaries of Gaussian dimension at least $2n - O(t)$ in time $poly(n,2^t,1/\varepsilon)$. Indeed, this class subsumes unitaries prepared by at most $O(t)$ non-Gaussian gates but also includes several unitaries that require up to $2^{O(t)}$ non-Gaussian gates to construct.
In addition, we give a $poly(n,1/\varepsilon)$-time algorithm to distinguish whether an n-mode unitary is of Gaussian dimension at least k or ε-far from all such unitaries in Frobenius distance, promised that one is the case. Along the way, we prove structural results about near-Gaussian fermionic unitaries that are likely to be of independent interest.

Authors: Vishnu Iyer

Title: StoCQS: stochastic strategy for Ansatz tree construction in Krylov-based linear solver

Abstract: Quantum algorithms for solving linear systems of equations have seen significant development since their inception. A method uses a classical combination of quantum states (CQS) along with the so-called Ansatz tree structure to construct approximate solutions to Ax=b with provable guarantees brought about by the Krylov subspace. However, the algorithm may require to construct the entire Ansatz tree to achieve the convergence guarantee, resulting in less efficiency when scaling to large-scale quantum machine learning. In this work, we propose StoCQS, an efficient strategy for Ansatz tree construction to solve linear systems with a convergence guarantee aided by importance sampling techniques and stochastic gradient descent, potentially with a reduced number of states. Our algorithm thus promises improved feasibility of implementing a quantum linear systems solver on large-scale quantum tasks with provable theoretical guarantees.

Authors: Xiufan Li, Soumik Adhikary and Patrick Rebentrost

Title: On the Cost of Training Adversarially-Robust Quantum Models

Abstract: In this paper, we study the cost of training parametrised quantum circuits (PQCs) that are employed in variational quantum algorithms (VQAs), which depends on the number of circuit-evaluations (CE). Since each circuit-execution is monetarily and temporally expensive, we treat cost and CE on an equal footing, and focus on the following questions:

Questions: How many CE are required to train quantum models to an epsilon-stationary solution of the objective function? Since the cost of estimating the gradients scales linearly with the number of parameters, can gradient-free algorithms be employed to get similar a performance at a lower expense?

Contributions: Our contributions towards answering these questions are threefold:

(a) As the number of CE depends linearly on the Lipschitz-smoothness constant of the objective function, we provide tightened bounds on it. Our result also generalises to arbitrary loss functions besides expectation values.

(b) For a circuit with d optimisable parameters, each (stochastic) gradient-based iteration consumes 2d evaluations of the circuit, which hinders the scaling-up of circuits to accommodate more parameters. We adapt a gradient-free algorithm called the Stochastic Three Points method to the VQA-setting, and demonstrate both theoretically and numerically that it reduces the number of requisite CE (over stochastic gradient descent) by a factor of d. Particularly, in conjunction with (a), we see a sesquilinear reduction in the cost of training PQCs - from the order of d^3.5 to d^2.

(c) Finally, we provide theoretical and empirical evidence to suggest that quantum models are inherently adversarially robust. This is due to the adversarial loss being lower and upper bounded by scalar multiples of the non-adversarial loss. This obviates the usual (expensive) overheads of adversarial training that arise from the estimation of input-gradients, and augmentation of the training dataset with adversarial examples.

Authors: Sayantan Pramanik, Chaitanya Murti, Chiranjib Bhattacharyya and M Girish Chandra

Title: Tolerant Testing of Stabilizer States with Mixed State Inputs

Abstract: We study the problem of tolerant testing of stabilizer states. In particular, we give the first such algorithm that accepts mixed state inputs. Formally, given a mixed state ρ that either has fidelity at least ε1 with some stabilizer pure state or fidelity at most ε2 with all such states, where

$\varepsilon 2 \leq \varepsilon 1^{O(1)}$, our algorithm distinguishes the two cases with sample complexity poly($1/\varepsilon 1$) and time complexity $O(n * \text{poly}(1/\varepsilon 1))$.
Authors: Daniel Liang and Vishnu Iyer

Title: Quantum state-agnostic work extraction (almost) without dissipation

Abstract: We investigate work extraction protocols designed to transfer the maximum possible energy to a battery using sequential access to N copies of an unknown pure qubit state. The core challenge is designing interactions to optimally balance two competing goals: charging the battery optimally using the qubit in hand, and acquiring more information qubit by qubit to improve energy harvesting in subsequent rounds. Here, we leverage the exploration-exploitation trade-off in reinforcement learning to develop adaptive strategies that achieve energy dissipation that scales only polylogarithmically in N. This represents an exponential improvement over current protocols based on full state tomography.
Authors: Josep Lumbreras, Ruo Cheng Huang, Yanglin Hu, Mile Gu and Marco Tomamichel

Title: Beyond Penrose tensor diagrams with the ZX-calculus: Applications to quantum computing, QML, condensed matter physics, and quantum gravity

Abstract: Since its inception, the ZX calculus has provided a diagrammatic way to reason about qubit quantum systems, representing any quantum computation or linear map between qubits as ZX-diagrams, a type of tensor network. It is recognized as a universal, sound, and complete language for qubit linear algebra. The study of quantum mechanics is deeply integrated with the representation theory of groups and algebras, especially SU(2) representation theory, essential for understanding quantum properties like spin and angular momentum. Despite advancements, a gap remained between diagrammatic languages like the ZX calculus and the algebraic structures in SU(2) representation theory which are key in fields such as quantum chemistry and condensed matter physics. The Penrose spin calculus was introduced to bridge this gap, integrating the ZX calculus's diagrammatic approach with SU(2)'s algebraic depth. Inspired by Penrose's work on spin networks and quantum geometry, it extends these concepts into a comprehensive diagrammatic language for SU(2) calculations, incorporating symmetric projectors and diagrammatic expressions of 3jm Wigner symbols, crucial for studying spin coupling and angular momentum. This facilitates an intuitive understanding and manipulation of these concepts, making SU(2) concepts more accessible and broadening their application in quantum computing and information theory.
Authors: Richard East, Quanlong Wang, Razin Shaikh, Lia Yeh, Boldizsár Poór and Bob Coecke

Title: Auxiliary-Free Replica Shadows: Efficient Estimation of Multiple Nonlinear Quantum Properties

Abstract: Efficiently measuring nonlinear properties is a significant yet challenging task from quantum information processing to many-body physics. Current methodologies often suffer from an exponential sampling cost or require auxiliary qubits and deep quantum circuits. To address these limitations, we propose an efficient auxiliary-free replica shadow (AFRS) framework, which leverages the power of the joint entangling operation on a few input replicas while integrating the mindset of shadow estimation. We rigorously prove that AFRS can offer exponential improvements in estimation accuracy compared with the conventional shadow method, and facilitate the simultaneous estimation of various nonlinear properties, unlike the destructive swap

test. Additionally, we introduce an advanced local-AFRS variant tailored to estimating local observables with constant-depth quantum circuits, significantly simplifying the experimental implementation. Our work paves the way for efficient and practical quantum measurements on near-term quantum hardware.
Authors: Qing Liu, Zihao Li, Xiao Yuan, Huangjun Zhu and You Zhou

Title: Quantum Advantage in Learning Quantum Dynamics

Abstract: One of the key challenges in quantum machine learning is finding relevant machine learning tasks with a provable quantum advantage. A natural candidate for this is learning unknown Hamiltonian dynamics. Here, we tackle the supervised learning version of this problem, where we are given random examples of the inputs to the dynamics as classical data, paired with the expectation values of some observable after the time evolution, as corresponding labels (outputs). The task is to replicate the corresponding input-output function. We prove that this task can yield provable exponential classical-quantum learning advantages under common complexity assumptions in natural settings. To design our quantum learning algorithms, we introduce a new method, which we call Fourier coefficient sampling for parametrized circuit functions, and which may be of independent interest. Furthermore, we discuss the limitations of generalizing our method to arbitrary quantum dynamics while maintaining provable guarantees. We explain that significant generalizations are impossible under certain complexity-theoretic assumptions, but nonetheless, we provide a heuristic kernel method, where we trade-off provable correctness for broader applicability.
Authors: Alice Barthe, Mahtab Yaghubi, Michele Grossi and Vedran Dunjko

Title: Quantum Recurrent Embedding Neural Network

Abstract: Quantum neural networks (QNNs) hold promise for leveraging quantum systems to tackle complex learning tasks, but their scalability has been hindered by severe trainability issues such as barren plateaus. In this work, we introduce the Quantum Recurrent Embedding Neural Network (QRENN), a novel quantum architecture inspired by fast-track information pathways in ResNets and grounded in general quantum circuit design principles. By employing tools from dynamical Lie algebra theory, we rigorously prove that QRENN circuits are trainable and can avoid barren plateaus. We validate the practical power of QRENN by applying it to two challenging quantum supervised learning tasks: classifying quantum Hamiltonians and detecting symmetry-protected topological (SPT) phases. Our model demonstrates high accuracy and robustness in both settings, showcasing its ability to learn nontrivial quantum features from data, suggesting a promising path toward scalable and reliable quantum machine learning models.
Authors: Mingrui Jing, Erdong Huang, Xiao Shi, Shengyu Zhang and Xin Wang

Title: Quantum Circuit simulation with a local time-dependent variational principle

Abstract: We introduce a new tensor network method for simulating quantum circuits based on a locally adaptive formulation of the Time-Dependent Variational Principle (TDVP), aimed at overcoming limitations of the widely used Time-Evolving Block Decimation (TEBD) algorithm. TEBD, though effective for short-range interactions and low entanglement growth, suffers from cumulative truncation errors and inefficiencies when simulating long-range gates, which require SWAP insertions that artificially increase entanglement and computational cost. The proposed method reformulates circuit simulation in the Schrödinger picture by interpreting each gate as a

small time-evolution step and projecting its generator onto the tangent space of the MPS manifold. This local dynamic TDVP approach enables more accurate simulations by preserving the MPS geometry and avoiding unnecessary truncations. It also supports dynamic bond dimension growth and direct simulation of long-range gates without introducing SWAP gates, thus maintaining computational efficiency. To implement this, single-qubit gates are directly contracted into MPS tensors, while multi-qubit gates are handled by slicing the MPS across affected regions and applying localized TDVP projections. A projector-splitting formalism is used to apply only the relevant summands for the affected qubits, significantly reducing computational overhead. For a gate acting on q qubits, the number of projectors is reduced from $2L-1$ (full TDVP) to $2q-1$, without loss of accuracy. Numerical benchmarks demonstrate that local dynamic TDVP matches TEBD in accuracy while requiring significantly fewer MPS parameters. A key example is a 36-qubit Trotterized 1D periodic Ising circuit with long-range interactions, where TEBD produces rapid entanglement growth and central bond dimension blow-up, while TDVP maintains way smoother, lower bond growth. These results suggest that local dynamic TDVP is a powerful and scalable alternative for simulating large quantum circuits, particularly those involving long-range gates or requiring dynamic entanglement handling.
Authors: Maximilian Fröhlich

Title: Variational LOCC-assisted quantum circuits for long-range entangled states

Abstract: Long-range entanglement is an important quantum resource, particularly for topological orders and quantum error correction. In reality, preparing long-range entangled states requires a deep unitary circuit, which poses significant experimental challenges. A promising avenue is offered by replacing some quantum resources with local operations and classical communication (LOCC). With these classical components, one can communicate outcomes of mid-circuit measurements in distant subsystems, substantially reducing circuit depth in many important cases. However, to prepare general long-range entangled states, finding LOCC-assisted circuits of a short depth remains an open question. Here, to address this challenge, we propose a quantum-classical hybrid algorithm to find optimal LOCC protocols for preparing ground states of given Hamiltonians. In our algorithm, we introduce an efficient way to estimate parameter gradients and use such gradients for variational optimization. Theoretically, we establish the conditions for the absence of barren plateaus, ensuring trainability at a large system size. Numerically, the algorithm accurately solves the ground state of long-range entangled models, such as the perturbed Greenberger–Horne–Zeilinger state and surface code. Our results demonstrate the advantage of our method over conventional unitary variational circuits: the practical advantage in the accuracy of estimated ground state energy and the theoretical advantage in creating long-range entanglement. This work has been published in Physical Review Letters [Phys. Rev. Lett. 134, 170601 (2025)].
Authors: Yuxuan Yan, Muzhou Ma, You Zhou and Xiongfeng Ma

Title: Machine Learning Derived Entanglement Witnesses with Trainable Measurements

Abstract: We propose a machine learning based approach to generating entanglement witnesses where the number of measurement settings may be directly specified. While previous methods of constructing witnesses express results in terms of a standard basis for local measurements like the Pauli basis, this algorithm directly trains and produces the specified number of Hermitian matrices, which can be represented by the same number of measurement settings on a given system. For \(N\) qudits of dimension \(d\) we use the fully separable eigenstates of the generators of \(SU(d)\) for each qudit as training data to determine the correct measurement settings, and we adjust the bias term of the witness with a new differential program to ensure maximal noise

tolerance and perfect accuracy. Additionally, we implemented adversarial training to produce witnesses of higher noise tolerance, requiring even fewer measurement settings. We provide an automated script to implement the above process that conveniently finds witnesses with better noise tolerance and/or fewer measurement settings than all existing methods in every nontrivial case we test. We apply this method to Bell states, GHZ states, W states, hypergraph states, and a range of qudit states. We consider systems from 2 - 5 qubits, bipartite qudits up to d = 10, and tripartite qutrits. Finally, we numerically verify each of the witnesses we generated with small test sets.
Authors: Aiden Rosebush, Alexander Greenwood and Li Qian

Title: Testing classical properties from quantum data
Abstract: Testing properties of Boolean functions is often dramatically faster than learning. However, this advantage usually disappears when testers are limited to random samples of the function---a natural setting for data science---rather than adaptive queries. In this work we investigate the \emph{quantum} version of this "data science scenario": quantum algorithms that test properties of a function f solely from quantum data in the form of copies of the function state $\ket{f}\propto \sum_x\ket{x,f(x)}$.
\textbf{New tests.} For three well-established properties---monotonicity, symmetry, and triangle-freeness---we show that the speedup lost when restricting classical testers to sampled data can be recovered by considering quantum data.
\textbf{Inadequacy of Fourier sampling.} Our new testers use techniques beyond quantum Fourier sampling, and we show that this is necessary. In particular, there is no constant-complexity tester for symmetry relying solely on Fourier sampling and random classical samples.
\textbf{Classical queries vs. quantum data.} We exhibit a testing problem that can be solved from $\mathcal{O}(1)$ classical queries but that requires $\Omega(2^{n/2})$ function state copies. The \textsc{Forrelation} problem provides a separation of the same magnitude in the opposite direction, so we conclude that quantum data and classical queries are "maximally incomparable" resources for testing.
\textbf{Towards lower bounds.} We also begin the study of \emph{lower bounds} for testing from quantum data. For quantum monotonicity testing, we prove that the ensembles used to prove exponential lower bounds for classical sample-based testing, do not yield any nontrivial lower bounds for testing from quantum data. New insights specific to quantum data will be required for proving copy complexity lower bounds for testing in this model.
Authors: Matthias C. Caro, Preksha Naik and Joseph Slote

Title: Quantum thermodynamics, semidefinite optimization, and quantum Boltzmann machines
Abstract: In quantum thermodynamics, a system is described by a Hamiltonian and a list of non-commuting charges representing conserved quantities like particle number or electric charge, and an important goal is to determine the system’s minimum energy in the presence of these conserved charges. Relevant quantum states of these systems are parameterized thermal states, alternatively known as quantum Boltzmann machines. In optimization theory, a semi-definite program (SDP) involves a linear objective function optimized over the cone of positive semi-definite operators intersected with an affine space. These problems arise from differing motivations in the physics and optimization communities and are phrased using very different terminology, yet they are essentially identical mathematically. By adopting Jaynes’ mindset motivated by quantum thermodynamics, we observe that minimizing free energy in the aforementioned thermodynamics problem, instead of energy, leads to an elegant solution in terms

of a dual chemical potential maximization problem that is concave in the chemical potential parameters. As such, one can employ standard (stochastic) gradient ascent methods to find the optimal values of these parameters, and these methods are guaranteed to converge quickly. At low temperature, the minimum free energy provides an excellent approximation for the minimum energy. We then show how this Jaynes-inspired gradient-ascent approach can be used in both first- and second-order classical and hybrid quantum–classical algorithms for minimizing energy, and equivalently, how it can be used for solving SDPs, with guarantees on the runtimes of the algorithms. The hybrid quantum–classical algorithms can be considered quantum
Boltzmann machine learning algorithms for energy minimization. The approach discussed here is well grounded in quantum thermodynamics and, as such, provides physical motivation underpinning why algorithms published fifty years after Jaynes' seminal work, including the matrix multiplicative weights update method, the matrix exponentiated gradient update method, and their quantum algorithmic generalizations, perform well at solving SDPs.
Authors: Nana Liu, Michele Minverini, Dhrumil Patel and Mark Wilde

Title: Optimal Haar random fermionic linear optics circuits

Abstract: Sampling unitary Fermionic Linear Optics (FLO), or matchgate circuits, has become a fundamental tool in quantum information. Such capability enables a large number of applications ranging from randomized benchmarking of continuous gate sets, to fermionic classical shadows. In this work, we introduce optimal algorithms to sample over the non-particle-preserving (active) and particle-preserving (passive) FLO Haar measures. In particular, we provide appropriate distributions for the gates of n-qubit parametrized circuits which produce random active and passive FLO. In contrast to previous approaches, which either incur classical $\mathcal{O}(n^3)$ compilation costs or have suboptimal depths, our methods directly output circuits which \textit{simultaneously} achieve an optimal down-to-the-constant-factor $\Theta(n)$ depth and
$\Theta(n^2)$ gate count; with only a $\Theta(n^2)$ classical overhead. Finally, we also provide quantum circuits to sample Clifford FLO with an optimal $\Theta(n^2)$ gate count.
Authors: Paolo Braccia, N. L. Diaz, Martin Larocca, M. Cerezo and Diego García-Martín

Title: The abelian state hidden subgroup problem: Learning stabilizer groups and beyond

Abstract: Identifying the symmetry properties of quantum states is a central theme in quantum information theory and quantum many-body physics. In this work, we investigate quantum
learning problems in which the goal is to identify a hidden symmetry of an unknown quantum state. Building on the recent formulation of the state hidden subgroup problem (StateHSP), we focus on abelian groups and develop a quantum algorithmic approach to efficiently learn any
hidden symmetry subgroup in this case. We showcase the versatility of the approach in three concrete applications: These are learning (i) qubit and qudit stabilizer groups, (ii) cuts along which a state is unentangled, and (iii) hidden translation symmetries. Taken together, our results underscore the potential of the StateHSP framework as a foundation for broader symmetry-based quantum learning tasks.
Authors: Marcel Hinsche, Jens Eisert and Jose Carrasco

Title: Scalable, hardware-efficient and noise-aware multivariate quantum state preparation

Abstract: Quantum state preparation of high-dimensional functions is an important component in many quantum algorithms. For these algorithms to provide a quantum advantage, both the classical and quantum subroutines required for state preparation need to be as efficient as

possible. Here we present classical algorithms based on tensor network methods that are efficient with regard to dimensionality. To avoid the barren plateau problem during the optimization, we devise a procedure that smoothly transforms the circuit from an easy-to-prepare initial function into the target multivariate function. We illustrate the approach by numerically optimising quantum
circuits for multivariate Gaussians of sizes up to 17 dimensions and using a total of 102 qubits. Additionally, we fine-tune the circuits composed of hardware-native quantum gates, taking realistic experimental noise into account, and demonstrate the experimental feasibility on
Quantinuum's H2 quantum computer, where we prepare a 9-dimensional Gaussian with polynomially decaying correlations using 54 qubits.
Authors: Marco Ballarin, Juan José García-Ripoll, David Hayes and Michael Lubasch

Title: Accelerating Inference for Multilayer Convolutional Neural Networks with Quantum Computers

Abstract: Fault-tolerant Quantum Processing Units (QPUs) promise to deliver exponential speed-ups in select computational tasks, but a clear recipe for integrating them into
existing classical deep-learning pipelines is missing. In this work, we focus on the setting of accelerating inference, where a pre-trained network outputs a probability distribution over possible outputs (e.g., classes or tokens), and ask: for which architectures, and under which Quantum Random Access Memory (QRAM) assumptions, can a fault-tolerant QPU asymptotically
outperform classical hardware? All the networks we consider are composed of fundamental residual blocks which consist of regularized multi-filter 2D convolutions, sigmoid activations, skip-connections, and layer normalisations, echoing the structure of ResNet.
To give end-to-end complexity bounds taking into account both input and memory assumptions, we consider three regimes (where the first two assume that QRAM is feasible). In regime 1, both the input tensor and network weights are given by QRAM. We prove that a network with an N-dimensional vectorized input, d residual block layers, and a final residual-linear-pooling layer can be implemented with $O(\text{polylog}(N)^d)$ inference cost, and suggest that known dequantization methods do not apply. Examples of tasks covered by this setting include those with repeated calls to a slowly-changing input, such as is the case with multi-turn LLM chat. Regime 2 keeps only the weights in QRAM; the input must be loaded at $O(N)$ cost. Even so, shallow multi-layer bilinear-style networks can achieve quartic (or better) speedups. Regime 3 has no QRAM, and so quantum speedup is unlikely in terms of either the dimension of the input, or the number of network parameters, and must instead exploit high dimensional latent feature transforms.
Authors: Arthur Rattew, Po-Wei Huang, Naixu Guo, Lirandë Pira and Patrick Rebentrost

QTML Tutorials

Title: Testing and Verification for Quantum Learning by Matthias Caro, University of Warwick

Abstract: In this tutorial, we will explore how techniques from property testing and interactive proofs can be leveraged in quantum learning theory. Concretely, we will consider the following questions: How can we test the assumptions of quantum learning algorithms? How can we delegate quantum learning algorithms in a verifiable manner? And how can we certify the hypotheses produced by quantum learning algorithms?

Title: Modern techniques in quantum algorithms by Zane Rossi, University of Tokyo

Abstract: Methods in the design and analysis of quantum algorithms have changed dramatically over the past ten years. These changes are rooted in new techniques for efficiently manipulating linear operators encoded in quantum computations: quantum signal processing (QSP) and quantum singular value transformation (QSVT). These algorithms, built from simple alternating circuit ansätze, enable one to apply functions to the spectra of linear operators encoded as sub-blocks of unitary matrices. This ability has proven surprisingly flexible—unifying, simplifying, and improving most quantum algorithms.

In this tutorial we review the construction and key properties of these algorithms, discuss how they are applied to solve disparate problems, and motivate recent extensions and improvements. Additionally, we highlight limitations of quantum algorithms for spectral mapping, and compare QSP/QSVT with quantum algorithms for manipulating linear operators in other ways (e.g., linear combination of Hamiltonian simulation). We aim to keep the tutorial accessible to both quantum and classical audiences, toward clarifying open questions in the intersection of quantum algorithms, functional analysis, and numerical linear algebra.

Title: Pauli Propagation Methods for Classical Simulation and Beyond by Armando Angrisani, EPFL

Abstract: As quantum devices grow up in size while remaining affected by noise, a central question is when and how they can surpass classical methods in practice. Pauli propagation has recently emerged as a powerful classical simulation framework, significantly raising the benchmark for demonstrating quantum advantage. By approximating the evolution of quantum operators in the Pauli basis, these methods achieve rigorous performance guarantees across a wide range of noisy and noiseless circuits. In this talk, I will present recent theoretical advances in Pauli propagation and highlight its potential well beyond classical simulation, including its use in quantum-inspired machine learning models and hybrid strategies that facilitate the practical deployment of near-term quantum devices. I will also discuss how the propagation framework can be generalized from discrete-variable to continuous-variable quantum systems.

QTML Industry Session

Title: Scalable quantum machine learning models in Fourier space by Joseph Bowles, Xanadu

Abstract: I will show how Fourier analysis can be used to construct quantum machine learning models that are massively scalable, while simultaneously encoding an important bias present in nearly all datasets. More specifically, we will see how using the quantum Fourier transform to build models in Fourier space leads to a class of universal generative machine learning models that can be trained entirely on classical hardware, enabling scaling to circuits with thousands of qubits and millions of parameters. By empirically implementing this training algorithm on large, real-world datasets, we will see how such models can be successfully trained despite the provable existence of barren plateaus under random parameter initialization. I will also introduce a new generative machine learning algorithm, called generative bandlimiting, that leverages this approach and exploits known biases in the Fourier spectra of data. Finally, I will conclude with some reflections on the overall direction of the field.

QTML Accepted Posters

Title: A competitive NISQ and qubit-efficient solver for the LABS problem

Abstract: Pauli Correlation Encoding (PCE) has recently been introduced as a qubit-efficient approach to combinatorial optimization problems within variational quantum algorithms (VQAs). The method offers a polynomial reduction in qubit count and a super-polynomial suppression of barren
plateaus. Moreover, it has been shown to feature a competitive performance with classical state-of-the-art (SOTA) methods on MaxCut. However, the latter problem class becomes truly hard only in the regime of very large input sizes. Here, we extend the PCE-based framework to solve the Low Autocorrelation Binary Sequences (LABS) problem. This is a notoriously hard problem with a single instance per problem size, considered a major benchmark for classical and quantum solvers. We simulate our PCE variational quantum solver for LABS instances of up to N = 44 binary variables using only n = 6 qubits and a brickwork circuit Ansatz of depth 10, with a total of 30 two-qubit gates,
i.e. well inside the NISQ regime. We observe a significant scaling advantage in the total time to (the exact) solution of our solver with respect to previous studies using the QAOA, and even a modest advantage with respect to the leading classical SOTA heuristic, given by Tabu search. Our findings point at PCE-based solvers as a promising quantum-inspired classical heuristic for hard-in-practice problems as well as a tool to reduce the resource requirements for actual quantum algorithms, with both fundamental and applied potential implications.

Authors: Marco Sciorilli, Lucas Borges, Giancarlo Camilo, Thiago O. Maciel, Askery Canabarro and Leandro Aolita

Title: A compositional framework for leveraging quantum learning to circumvent no-go theorems

Abstract: Quantum theory is subject to a variety of no-go theorems. Specifically, these are functions that cannot be implemented via quantum computations under the assumptions of exactness, determinism, and universality. Much work has been done to explore how these no-go theorems can be 'circumvented' if one of these is relaxed. This work focusses on the final one of these assumptions. Namely, how much about our input is it necessary to learn in order to implement a desired function?
Here we propose a framework for constructing and analysing implementations of functions on quantum objects which depend upon some learnt knowledge about the input in order to be implemented. We show that these implementations along with the requisite learnt knowledge can be composed to form a complete compositional framework that can be used to construct complicated functions and know exactly what must be first learnt in order to execute them successfully.

Authors: Tim Forrer, Matthew Wilson, Philip Taranto, Jisho Miyazaki and Mio Murao

Title: A faster converging qDRIFT algorithm with application to Hamiltonian data encoding

Abstract: Hamiltonian-based data encoding is a vital component in quantum machine learning, providing powerful means to map classical data into quantum states. However, the

implementations have been limited by the computational resources required for simulating complex Hamiltonians. The qDRIFT algorithm can be used for Hamiltonian data encoding however the gate complexity exhibits unfavourable quadratic scaling with respect to the sum of Hamiltonian coefficients. In this talk, I present a refined theoretical analysis of the qDRIFT algorithm, revealing a significantly
tighter error bound achieved by leveraging the integral error representation of Taylor's theorem in conjunction with Jensen's inequality. This new result eliminates the problematic quadratic dependence, establishing that the required number of qDRIFT steps now scales linearly with
respect to the sum of coefficients In Hamiltonian. This fundamental improvement makes Hamiltonian-based data encoding and other quantum simulation tasks more efficient. The results further generalise to simulating open quantum systems and quantum chemistry applications.
Authors: Ian Joel David, Ilya Sinayskiy and Francesco Petruccione

Title: A Practical Cross-Platform, Multi-Algorithm Study of Quantum Optimisation for Configurational Analysis of Materials

Abstract: Quantum computers show potential for achieving computational advantage over classical computers, with many candidate applications in the domains of chemistry and materials science. We consider the well-studied problem of configurational analysis of materials and, more specifically, finding the lowest energy configuration of defective graphene structures. This problem acts as a test-case which allows us to study various algorithms that are applicable to Quadratic Unconstrained Binary Optimisation (QUBO) problems, which are generally classically intractable exactly. To solve this problem, we implement two methods, the Variational Quantum Eigensolver (VQE) and Quantum Annealing (QA), on commercially-available gate-based and quantum annealing devices that are accessible via Quantum-Computing-as-a-Service (QCaaS) models. To analyse the performance of these algorithms, we use a toolbox of relevant metrics and compare performance against three classical algorithms. We employ quantum methods to solve fully-connected QUBOs of up to 72 variables, and find that algorithm performance beyond this is restricted by device connectivity, noise and classical computation time overheads. The applicability of our approach goes beyond the selected configurational analysis test-case, and we anticipate that our approach will be of use for optimisation problems in general. arXiv:2504.06885.
Authors: Kieran McDowall, Theodoros Kapourniotis, Christopher Oliver, Phalgun Lolur and Konstantinos Georgopoulos

Title: A Resource-Efficient Quantum-Classical Model for Protein–Ligand Binding Affinity Prediction

Abstract: We propose a resource-efficient hybrid quantum-classical model for protein-ligand affinity prediction, building on Mavrommati et al. (2024) by replacing variational quantum regressors (VQRs) with a semi-adaptive ansatz using angle encoding and layer-wise entanglement. Our method reduces
trainable parameters by up to 96.88%, gate count by up to 93.37%, and depth by up to 96.24%. Evaluated on three benchmark datasets, it is comparable to existing performance while improving noisy intermediate-scale quantum (NISQ) suitability.
Authors: Tara Kit, Leanghok Hour, Muyleang Ing and Youngsun Han

Title: A Study on Stabilizer R\'enyi Entropy Estimation using Machine Learning

Abstract: We propose a supervised Machine Learning approach to estimate the stabilizer Renyi entropy SRE, a measure of nonstabilizerness. The nonstabilizerness of a quantum state quantifies to what extent it diverges from the set of stabilizer states, which can be efficiently simulated on classical computers. Nonstabilizerness is thus a fundamental resource for quantum advantage. This paper focuses on the SRE because of its computational properties and suitability for experimental measures on quantum processors. However, estimating SRE of arbitrary quantum states is a computationally hard problem. In this paper, we frame SRE estimation as a regression task and train a Random Forest Regressor and a Support Vector Regressor on a comprehensive dataset, including both unstructured random quantum circuits and structured circuits derived from the 1D transverse Ising model TIM. We compare two input representations: one based on the classical shadow protocol, and another encoding circuit-level features. While classical shadows achieves lower training error, circuit-level features exhibit better generalization performance. Moreover, we assess the generalization capabilities of the models on out-of-distribution instances. The experimental results are random quantum circuits require more complex models and informative representations, on the TIM dataset the SVR achieves a mean squared error of 0.06 and 0.1, respectively, generalizing to unseen circuit depth and number of qubits.

Authors: Vincenzo Lipardi, Domenica Dibenedetto, Georgios Stamoulis and Mark H.M. Winands

Title: A Unified Frequency Principle for Quantum and Classical Machine Learning

Abstract: The frequency principle describes the tendency of machine learning models to capture low-frequency components more efficiently during training. This phenomenon has been widely observed in deep neural networks and more recently in quantum machine learning models. In this work we provide a unified and rigorous mathematical theorem that establishes the existence of the frequency principle in both classical and quantum machine learning frameworks. Our approach relies on a general theorem derived from the spectral integral representation of unbounded self-adjoint operators.

Authors: Rundi Lu, Ruiqi Zhang, Weikang Li, Dong-Ling Deng and Zhengwei Liu

Title: ADAPT-VQE with Operator Removal

Abstract: ADAPT-VQE is a well-known algorithm for near-term quantum computers that creates a state preparation circuit by selecting an operator from a predefined pool in each iteration, after which a full VQE optimization is completed. The selection criterion is based on energy derivatives, reflecting the expectation that operators with a high impact on the energy locally will also have a high impact once a full optimization is performed. However, this selection criterion is an imperfect heuristic—an operator with a higher magnitude energy derivative might result in a lower energy change. Ideally, we would add the operator corresponding to the highest energy change, but finding it would imply a prohibitive number of optimizations. In this work, we propose a cost-efficient modification of the ADAPT-VQE algorithm that partly compensates for the shortcomings of the gradient heuristics by removing poorly performing operators based on available information. Our protocol addresses three fundamental questions in an operator removal protocol: (1) how to decide which operators are good candidates for removal without incurring additional measurements, (2) how to prevent the removal of operators which are important to the ansatz, and (3) how to avoid re-adding removed operators too soon.

Authors: Mafalda Ramôa

Title: Advanced Ensemble Smart Classifications for Nifty Smart Market Trends

Abstract: For investors to make wise decisions in the ever-changing stock market, accurate and timely insights are essential. Predicting market trends is still difficult because stock prices are inherently volatile. Using a novel hybrid algorithm that combines prediction accuracy is increased using ensemble learning methods that use Random Forest, Gradient Boosting, and Long Short-Term Memory (LSTM) networks. We present an advanced ensemble-based smart classification approach for predicting NIFTY market trends in this paper. In order to capture intricate market patterns and sentiment-driven fluctuations, our model makes use of historical stock data, technical indicators, financial news, stock forum discussions, and sentiment analysis from social media. In addition to efficiently processing sequential data and detecting long-term dependencies, ensemble techniques like Gradient Boosting with Random Forest increase prediction precision and robustness. We carry out a thorough analysis of stock market prediction techniques, addressing significant issues. We evaluate performance metrics like accuracy, Mean Absolute Error (MAE), and Root Mean Square Error (RMSE). Our model outperforms traditional methods with an amazing 99.9% accuracy. We demonstrate how this work is intrinsically transformative by showing how the hybrid deep learning and ensemble frameworks significantly improved stock market prediction, giving investors more trustworthy and useful market information.

Authors: Mandadi Sindhu, M.Mohan Sai Reddy and Devireddy Sai Santhosh Reddy

Title: AiDE-Q: Synthetic Labeled Datasets Can Enhance Learning Models for Quantum Property Estimation

Abstract: Existing deep learning (DL)-based models for quantum property estimation (QPE) typically assume access to large-scale, noiseless labeled datasets generated by infinite sampling. This assumption raises fundamental concerns regarding their practical applicability. In this manuscript, we propose AiDE-Q (automatic data engine for quantum property estimation), a method designed to unlock the potential of DL-based models using limited quantum resources. AiDE-Q achieves this by iteratively generating high-quality synthetic labeled datasets. To validate the effectiveness of AiDE-Q, we perform extensive numerical simulations on a diverse set of quantum many-body and molecular systems, including those with up to 50 qubits. The results demonstrate that AiDE-Q significantly enhances the prediction performance of various reference learning models, with improvements reaching up to 14.2%. Our work paves the way for more efficient and practical applications of deep learning in QPE.

Authors: Xinbiao Wang, Yuxuan Du, Zihan Lou, Yang Qian, Kaining Zhang, Yong Luo, Bo Du and Dacheng Tao

Title: AI-Powered Noisy Quantum Emulation: Generalized Gate-Based Protocols for Hardware-Agnostic Simulation

Abstract: Quantum computer emulators model the behavior and error rates of specific quantum processors. Without accurate noise models in these emulators, it is challenging for users to optimize and debug executable quantum programs prior to running them on the quantum device, as device-specific noise is not properly accounted for. To overcome this challenge, we introduce a general protocol to approximate device-specific emulators without requiring pulse-level control. By applying machine learning to data obtained from gate set tomography, we construct a device-specific emulator by predicting the noise model input parameters that best match the target device. We demonstrate the effectiveness of our protocol's emulator in estimating the unitary coupled cluster energy of the H_2 molecule and compare the results with those from actual

quantum hardware. Remarkably, our noise model captures device noise with high accuracy, achieving a mean absolute difference of just 0.3% in expectation value relative to the state-vector simulation.
Authors: Matthew Ho, Jun Yong Khoo, Adrian Mak and Stefano Carrazza

Title: An Attention-Based Quantum Phase Transition Detection on NISQ Devices

Abstract: Learning many-body quantum states and identifying quantum phase transitions remain major challenges in quantum many-body physics. Classical machine learning methods offer promising tools to address these difficulties. In this work, we propose a novel framework that bypasses the need to measure physical observables by directly learning from the parameters of parameterized quantum circuits.

Our model incorporates an attention mechanism with an encoder-decoder architecture, enabling it to learn intricate correlations among circuit parameters. These encoded representations capture hidden correlations among circuit parameters based on the quantum phase information in an unsupervised manner, enabling them to naturally cluster according to different quantum phases. Moreover, the encoded form serves as a compact classical representation of parameterized quantum circuits and the corresponding many-body quantum states, allowing for efficient generation of representative quantum states associated with specific phases.

We apply our framework to a variety of many-body quantum systems, including symmetry-protected topological (SPT) phases, demonstrating its broad applicability and strong performance, particularly in identifying topological quantum phase transitions.

Authors: Li Xin and Zhang-Qi Yin

Title: Balancing Expressivity and Learnability in Quantum Kernel Bandit Optimization

Abstract: We investigate Gaussian Process (GP) bandit optimization utilizing quantum kernels. While quantum kernels enable embedding data into high-dimensional Hilbert spaces—potentially offering enhanced expressivity or a "quantum advantage"—this property can pose challenges in bandit learning. Specifically, employing full quantum kernels naïvely may lead to increased model complexity and, consequently, higher cumulative regret impacting the learnability.

To address this, we explore the use of projected quantum kernels and classical kernel approximation techniques, which effectively reduce feature dimensionality while preserving essential quantum properties. We demonstrate that these approaches can yield improved regret bounds by strategically balancing approximation error and information gain. Empirical results show that they significantly outperform models based on full quantum kernels in bandit optimization tasks.

Additionally, we analyze how to select the optimal model complexity to achieve a favorable trade-off between expressivity and learnability. Our methods also substantially reduce computational costs by simplifying kernel-based inference to linear models, since only a finite set of reduced features is required.

Authors: Yuqi Huang, Vincent Y. F. Tan and Sharu Jose

Title: Barren plateau-free and noise-robust quantum advantage for learning data with group symmetries

Abstract: Quantum machine learning (QML) can outperform classical machine learning for certain structured data problems. Recent work has demonstrated that data exhibiting a specific kind of group structure can be more efficiently learned with QML techniques than with their classical

counterparts. However, a major obstacle in QML is the phenomenon of barren plateaus, where gradients vanish exponentially and training becomes impossible.
In this work, we use kernel methods and prove that for this group-structured problem, learning is free of barren plateaus. This result is a significant step forward in understanding the trainability of QML methods.
Authors: Laura Henderson, Kerstin Beer, Salini Karuvade, Riddhi Gupta and Angela White

Title: Bayesian Learning of Quantum Hardware Dynamics

Abstract: Bayesian inference can be fruitfully applied to the characterization of key processes in quantum computing, with applications in sensing, phase estimation, device calibration, and more. However, its performance is limited by the quality of the numerical representation of Bayesian probability distributions. This limitation is particularly pronounced in challenging scenarios, such as learning the dynamics of open quantum systems, which may have a high number of parameters and often exhibit multi-modality due to redundant explanations.
In this work, we review the statistical methods used for Bayesian quantum learning, numerically analyze their performance, discuss their limitations, and propose more robust alternatives. We apply advanced statistical techniques to the characterization of open quantum systems via Bayesian inference to optimize the inference process and design appropriate control of the quantum systems. While Bayesian inference has often been applied to quantum problems, these cases typically target uni-dimensional estimation problems where simple statistical tools suffice. We demonstrate the shortcomings of these tools through numerical simulations and propose more robust alternatives. We also overview Bayesian experimental design and related topics, discussing how well-known heuristics or commonly adopted techniques can fail under certain circumstances, and propose and test alternatives.
In particular, we report and compare results from various techniques, including sequential importance resampling with the Liu-West filter and Markov Chain Monte Carlo (MCMC) kernels, Hamiltonian Monte Carlo (HMC), stochastic gradient HMC with and without friction, HMC with energy-conserving subsampling, random walk Metropolis-Hastings (RWM), tempered likelihood estimation, block pseudo-marginal Metropolis-Hastings with subsampling, hybrid approaches that adaptively switch between HMC and RWM, and Gaussian rejection filtering. We also propose generalized adaptive heuristics for multi-modal likelihoods.
As a benchmark, we apply these techniques to several quantum problems, including the estimation of decoherence effects (energy loss and dephasing) and frequencies in quantum hardware. We use the algorithms to calibrate IBMQ superconducting quantum hardware, observing improved performance compared to Qiskit's default fitters, especially in low-data regimes. This improvement can be further enhanced with subsampling strategies. For Hahn echo and Ramsey experiments, we achieve uncertainties respectively 3 and 10 times smaller than default methods using the same number of experiments. Conversely, we match Qiskit's performance while using up to 99.5% less data. We additionally test an adaptive heuristic for Hahn-Ramsey experiments, achieving an eightfold reduction in uncertainty for the same number of shots, and explore particularities of coherence time estimation that contradict typically adopted strategies in quantum characterization.
Our algorithms excel in problems where experimental data collection is costly, real-time estimation is required, or the likelihood functions are difficult to sample from. These findings have broad applications in challenging quantum characterization tasks, particularly in learning the dynamics of open quantum systems.
Authors: Alexandra Ramôa, Raffaele Santagati and Nathan Wiebe

Title: Benchmarking Quantum Algorithms for Gaussian Process Regression

Abstract: Gaussian Process Regression is a well-known machine learning technique for which several quantum algorithms have been proposed. We demonstrate that in a wide range of scenarios, these algorithms show no exponential speedup. We achieve this by rigorously proving that the condition number of a kernel matrix scales at least linearly with the matrix size under general assumptions on the data and kernel.

Additionally, we prove that the sparsity and Frobenius norm of a kernel matrix scale linearly under similar assumptions. The implications for the runtime of quantum algorithms are independent of the complexity of loading classical data onto a quantum computer and also apply to dequantized algorithms. We supplement our theoretical analysis with numerical verification for popular kernels in machine learning.

Authors: Dominic Lowe, Myungshik Kim and Roberto Bondesan

Title: Bridging Remote Sensing and Quantum Computing: Snow Depth Estimation with LSTM and QLSTM

Abstract: Accurate and timely estimation of snow depth is essential for effective hydrological forecasting, flood risk management, and climate modeling, especially in cold and remote regions where in-situ data collection is logistically challenging and sparse. This research proposes a comprehensive, data-driven approach to model daily snow depth by integrating multivariate environmental inputs using both classical and quantum deep learning methods. Specifically, we compare the performance of a Long Short-Term Memory (LSTM) neural network, and a Quantum Long Short-Term Memory (QLSTM) model trained on spaceborne Synthetic Aperture Radar (SAR) statistics and meteorological variables.

The dataset combines irregular SAR observations from the RADARSAT Constellation Mission (RCM) processed to extract backscatter intensity statistics (mean and standard deviation of HH and HV polarizations)—with daily meteorological variables, such as precipitation and temperature, and sparse in-situ snow depth measurements. Due to temporal inconsistencies in SAR acquisition (~1 image every four days) and snow depth sampling (typically every 15 days), a temporal harmonization strategy was

applied: SAR values were filled using backward fill interpolation, while snow depth was linearly interpolated to daily resolution. This fusion of multi-source data enables the construction of a high-resolution time series dataset suitable for sequential modeling.

The LSTM model serves as a classical baseline due to its proven ability to capture long-range temporal dependencies in time-series data. The QLSTM model, in contrast, integrates quantum computing principles through a hybrid architecture implemented in PennyLane and PyTorch. In QLSTM, classical input features are encoded into quantum states using a parameterized quantum circuit, which performs entangling operations across qubits to capture complex feature correlations. The quantum-transformed features are then passed through an LSTM cell, enabling temporal learning within a quantum-enhanced feature space. This setup leverages the expressive power of quantum states (superposition and entanglement) to potentially improve learning of non-linear, high-dimensional relationships inherent in snow dynamics.

Both models were trained and evaluated on the same dataset, with performance metrics including Root Mean Square Error (RMSE), Mean Absolute Error (MAE), R^2, and training time. We hypothesize that while the classical LSTM will deliver reliable and computationally efficient performance, the QLSTM may outperform it in generalization accuracy and sensitivity to subtle temporal variations—though at the cost of higher computational demands due to quantum circuit simulation.

This study contributes to emerging efforts to apply quantum machine learning in Earth system science and aims to provide practical insights into the scalability, feasibility, and performance trade-offs of using
quantum-enhanced models for environmental monitoring in data-limited contexts.
Authors: Charitha Pathipati, Tirupati Bolisetti, Girish Sankar, Sri Harsha Thota and Ram Balachandar

Title: Canonical Quantization of a Memristive Leaky Integrate-and-Fire Neuron Circuit

Abstract: We develop a quantized memristive Leaky Integrate-and-Fire (LIF) neuron as a unifying paradigm at the intersection of neuromorphic engineering and open quantum systems. In classical neuromorphic platforms, memristors—non-linear two-terminal elements whose resistance depends on historic charge flow—provide circuits with intrinsic memory and plasticity, closely mimicking biological neurons and synapses. By promoting both the membrane capacitance and the memristor to quantum degrees of freedom, we construct a fully Hamiltonian description within the framework of circuit quantum electrodynamics. This quantization preserves the hallmark features of the LIF model, such as voltage integration, leak, thresholding, and reset, while embedding them in a coherent, open-system setting amenable to quantum control techniques.
Starting from the classical memristive LIF circuit, we apply canonical quantization to each lumped element, yielding flux and charge operators that satisfy the usual commutation relations. Coupling the resulting quantum LC oscillator to a transmission-line continuum of bosonic modes through a coupling capacitor provides a microscopic origin for memristive dissipation. Tracing out the bath in the weak-coupling, Born–Markov limit leads to a GKSL master equation with a dissipator that generates the leak dynamics. This approach offers a first-principles derivation of memristive and spiking behavior fully embedded in quantum mechanics.
We validate the model through numerical simulations in the weak-coupling, adiabatic regime. The quantized memristor reproduces the characteristic pinched hysteresis loops in the current–voltage plane, and the stochastic jump process yields well-defined spike trains with refractory intervals. Importantly, the memristance operator evolves in time according to the history of quantum charge, demonstrating genuine quantum memory effects while converging to classical dynamics in the appropriate limit. These results confirm that our model simultaneously captures quantum coherence phenomena and the essential features of neuronal information processing.
By establishing this quantized memristive LIF neuron, we lay the groundwork for quantum spiking neural networks, which can leverage quantum superposition and entanglement to encode, store, and process information in ways unattainable with classical hardware. In particular, such networks promise novel quantum machine learning protocols in which synaptic weights and neuronal thresholds are represented by quantum states and operators, enabling exponentially large Hilbert-space embeddings and potentially faster training and inference. Furthermore, our framework invites exploration of regimes beyond the Born–Markov approximation, where non-Markovian dynamics and genuinely quantum spiking behaviors, such as superposed firing trajectories, may emerge. This work opens new avenues for biologically inspired, quantum-native computing paradigms.
Authors: Dean Brand, Domenica Dibenedetto and Francesco Petruccione

Title: Certifying Adversarial Robustness in Quantum Machine Learning: From Theory to Physical Validation

Abstract: As with classical neural networks, quantum machine learning (QML) models are vulnerable to small input perturbations that can significantly alter output predictions. Ensuring the

robustness of QML models, particularly on NISQ hardware, is therefore a fundamental step toward trustworthy quantum AI.
This paper presents a comprehensive framework for certifying adversarial robustness in QML. Our core contribution is a fidelity-based robustness bound, computable directly from the measurement outcome distribution, which enables both formal certification and empirical estimation on real quantum devices. Additionally, the optimal bound can be computed via semidefinite programming (SDP) with full knowledge of the quantum machine learning models.
We incorporate these results into (1) an efficient formal verification framework; (2) VeriQR, the first dedicated QML robustness certification tool; and (3) the first experimental benchmark of quantum adversarial robustness on a 20-qubit superconducting processor. Together, these systematic advances enable scalable, physically grounded robustness evaluation of QML.
Authors: Ji Guan

Title: Certifying Optimality of VQA Solutions via Sparse SOS Hierarchies
Abstract: We propose a certification framework for Variational Quantum Algorithms (VQAs) based on sparse Sum-of-Squares (SOS) relaxations. For a broad class of ansätze, including Quantum Approximation Optimization Algorithms (QAOA), we show that the VQA cost function admits a sparse Hermitian trigonometric polynomial representation. This structure enables the construction of sparse SOS hierarchies that provide two-sided bounds on the suboptimality of any candidate parameter of the VQA.
The error in the suboptimality gap at the d-th level of the hierarchy scales as $O(1/d^2)$. To our knowledge, this provides the first certified suboptimality guarantees for VQAs across a broad class of ansätze and offers a tractable post hoc validation method for VQA solutions.
Authors: Georgios Korpas, Wayne Lin, Iosif Sakos and Antonios Varvitsiotis

Title: Challenges and limitations of quantum kernel methods
Abstract: This contribution provides comprehensive insights into the key findings of our recent works on benchmarking quantum kernel methods (QKMs) and elucidates the role of bandwidth tuning, focusing on its impact on generalization and classical tractability. Overall, our results highlight critical aspects of QKMs and emphasize the need to approach quantum machine learning research from two perspectives: identifying datasets that require leveraging quantum-specific properties and exploring how corresponding model designs should be structured.
Authors: Jan Schnabel, Roberto Flórez-Ablan and Marco Roth

Title: Characterizing quantum resourcefulness via group-Fourier decompositions
Abstract: In this work, we present a general framework for studying the resourcefulness of pure states in quantum resource theories (QRTs) whose free operations arise from the unitary representation of a group. We argue that the group Fourier decompositions (GFDs) of a state—its projection onto the irreducible representations (irreps) of the Hilbert space, operator space, and their tensor products—constitute fingerprints of resourcefulness and complexity.
By focusing on the norm of the irrep projections, dubbed GFD purities, we find that low-resource states occupy the small-dimensional irreps of operator space, whereas high-resource states have support in more, and higher-dimensional, irreps. This behavior not only resembles phenomena observed in classical harmonic analysis but is also universal across QRTs such as entanglement, fermionic Gaussianity, SU(2), and Clifford stabilizerness.

Finally, we show that GFD purities carry operational meaning: they can serve as resourcefulness witnesses, provide notions of state compressibility, and capture features that are inequivalent to those described by a state's extent—its minimal decomposition into free states.
Authors: Paolo Braccia, Pablo Bermejo, Antonio Anna Mele, Nahuel Diaz, Andrew Deneris, Martin Larocca and Marco Cerezo

Title: Classical-quantum hybrid support vector data description for one-class classification

Abstract: One-class classification (OCC) is a crucial task in machine learning, with applications such as anomaly detection and quality control. As modern datasets grow in complexity and dimensionality, there is an increasing need for advanced OCC techniques.
We propose Neural Quantum Support Vector Data Description (NQSVDD), a classical-quantum hybrid algorithm tailored for OCC. Our approach first learns a feature representation via a neural network combined with a quantum feature map, embedding the input data into a high-dimensional Hilbert space. A variational quantum circuit is then used to project the embedded features into a latent space, where the optimization process identifies the minimum-volume enclosing hypersphere as the decision boundary.
Experiments on the MNIST and Fashion-MNIST datasets show that NQSVDD outperforms conventional baselines in both accuracy and parameter efficiency.
Authors: Changjae Im, Hyeondo Oh and Daniel K. Park

Title: Classification of Quantum Correlations via Quantum-inspired Machine Learning

Abstract: We introduce a quantum-inspired classification framework based on the Pretty Good Measurement (PGM), aimed at identifying and discriminating quantum states according to their correlation structure. Specifically, our model distinguishes among factorized, separable, and entangled states, and determines the presence or absence of non-local correlations. Building on prior work limited to pure states, we generalize the PGM-based approach to encompass both pure and mixed quantum ensembles, with applications to systems ranging from two to five qubits.
The classifier encodes the training data into a measurement scheme derived from convex combinations of representative states and assigns labels based on maximum likelihood inference using PGM-defined POVMs. We benchmark our method against a suite of classical machine learning algorithms, demonstrating superior performance in balanced accuracy, particularly in detecting entanglement and factorization.
Moreover, we introduce a binary classification scheme for non-locality based on the CHSH and Svetlichny inequalities for two- and three-qubit systems, respectively, capturing higher-order quantum correlations in a data-driven manner. The results highlight not only the scalability, but also the physical interpretability and robustness of the PGM Classifier in noisy and mixed-state scenarios.
By incorporating key operational features of quantum mechanics into a supervised learning framework, this approach provides a principled and computationally efficient tool for quantum state discrimination. Our findings indicate that the PGM-based classification scheme may serve as a promising component within hybrid quantum-classical workflows, particularly for tasks involving the analysis of quantum correlations and the characterization of quantum resources.
Authors: Giuseppe Sergioli, Roberto Giuntini, Andres Camillo Granda Arango, Carlo Cuccu and Carla Sophie Rieger

Title: Conservative Quantum Offline Model-Based Optimization

Abstract: Offline model-based optimization (MBO) involves optimizing a black-box objective function using only a static dataset of prior evaluations, without performing any new queries. Quantum extremal learning (QEL) is a recent quantum machine learning approach that trains a parametrized quantum circuit as a surrogate model of the objective function and then uses its gradients to propose a candidate maximizer.

However, like classical surrogate models, QEL can suffer from overly optimistic predictions on inputs outside the support of the data, potentially leading to poor solutions. In this paper, we propose integrating QEL with conservative objective models (COM), a regularization technique that encourages cautious predictions on out-of-distribution inputs. The resulting hybrid algorithm, COM-QEL, leverages the expressive power of quantum neural surrogates while safeguarding generalization through prudent conservative modeling.

Authors: Kristian Sotirov, Annie E. Paine, Savvas Varsamopoulos, Antonio A. Gentile and Osvaldo Simeone

Title: Data Clustering as a Quantum Computing Use-Cas

Abstract: The research projects QORA and QORA II, funded by the Ministry of Economics, Labor and Tourism Baden-Württemberg, explored the potential applications of quantum computing (QC) with resilient algorithms over the past four years. In particular, use cases in the domain of finance were investigated, including portfolio optimization and feature selection for a credit scoring algorithm. To assess business impact, potential quantum methods were compared to established classical (non-quantum) computing approaches.

Authors: Prof. Dr. Gerhard Hellstern

Title: Deep Reinforcement Learning for real-time context-aware gate calibration

Abstract: Quantum computing faces challenges in reducing error rates below necessary thresholds, partly because traditional calibration methods do not account for the specific contexts in which quantum gates operate, each of which typically carries a unique noise signature.

We present a new framework that leverages model-free reinforcement learning (RL) to dynamically calibrate quantum gates based on their circuit context, adding an additional layer of error suppression during circuit execution.

Our work highlights two main contributions:

Simulation results: We suppress contextual coherent noise, such as classical microwave crosstalk in superconducting qubits. By training an RL agent to optimize two-qubit gates within noisy circuit layers, we achieve circuit fidelities exceeding 99.99%.

Experimental integration: We embed the RL workflow directly on the control system to minimize communication and compilation latency while maintaining flexible parametrization of custom calibrations. This enables contextual calibration within minutes of wall-clock time, significantly faster than traditional closed-loop methods, which often require many hours.

Additionally, our method does not require extra controller memory or latency, as it uses real-time inference from a classical neural network running asynchronously with the quantum circuit. This adaptability supports both fast agent training and fault-tolerant circuit executions, providing an additional layer of noise robustness at no extra cost. We also demonstrate that the agent can generalize to untrained circuit contexts by carefully selecting the training dataset.

By combining circuit-level considerations with pulse-level optimization, our framework offers a robust strategy for achieving and maintaining low physical error rates, advancing scalable quantum computing. This approach, together with the cross-layer abstraction framework we created,

represents one of the first steps toward low-latency hybrid classical-quantum computing, with applications in real-time error correction and decoding becoming increasingly important.
Authors: Arthur Strauss, Lukas Voss, Aniket Chatterjee and Hui Khoon Ng

Title: Dequantization and expressivity in photonic quantum Fourier models

Abstract: In this work, we study the models emerging from linear optical circuits and their augmented versions with non-linearity, such as feedforward adaptivity or state injection. These architectures are promising for near-term applications of quantum computing in learning tasks.
Meanwhile, Fourier representations of quantum circuits have been introduced as a powerful tool for analyzing variational algorithms (VQAs). Recent works have shown that this approach can inform the design of surrogate models and reveal separation results in the context of learning tasks.
Authors: Hugo Thomas, Hela Mhiri, Leo Monbroussou, Zoë Holmes and Elham Kashefi

Title: Design nearly optimal quantum algorithm for linear differential equations via Lindbladians

Abstract: Solving linear ordinary differential equations (ODEs) is a promising application for quantum computers to demonstrate exponential advantages. A key challenge in designing quantum ODE algorithms is embedding non-unitary dynamics into inherently unitary quantum circuits.
In this work, we propose a new quantum algorithm for solving ODEs by leveraging open quantum systems. Specifically, we introduce a technique called non-diagonal density matrix encoding, which utilizes the inherent non-unitary dynamics of Lindbladians to encode general linear ODEs into the non-diagonal blocks of density matrices. This framework allows for the design of quantum algorithms that are both theoretically simple and high performing.
Combined with state-of-the-art quantum Lindbladian simulation algorithms, our approach can outperform existing quantum ODE algorithms and achieve near-optimal dependence on all parameters under a plausible input model. Applications of our algorithm include Gibbs state preparation and partition function estimation.
Authors: Zhong-Xia Shang, Naixu Guo, Dong An and Qi Zhao

Title: Designing Privacy-Preserving Architectures in Quantum Federated Learning

Abstract: Quantum Federated Learning (QFL) provides a promising approach for secure and privacy-preserving quantum machine learning by enabling distributed training across multiple quantum or hybrid clients without sharing raw data. This emerging field combines quantum computing with federated learning to support intelligent systems that can learn from distributed, sensitive datasets while respecting privacy constraints and limited quantum resources.
In this work, we explore a range of QFL techniques that integrate quantum circuits with various learning strategies to address key challenges, including data heterogeneity, scalability, and the vulnerability of quantum systems to noise. These approaches leverage quantum spiking dynamics, differential privacy mechanisms, and encrypted model updates via homomorphic encryption to maintain model performance while safeguarding information.
The techniques are evaluated on diverse datasets, including structured and multimodal ones, to assess their effectiveness in real-world, privacy-critical scenarios. Beyond performance, the architectures demonstrate flexibility in adapting to different quantum backends, including simulators and real quantum hardware.

Together, these methods contribute to the design of scalable, privacy-aware quantum intelligence that supports collaborative learning in domains such as healthcare diagnostics, genomic analysis, and secure financial systems. This work highlights the practical potential of QFL and lays the groundwork for future research focused on robustness, interoperability, and deployment in noisy intermediate-scale quantum (NISQ) environments.
Authors: Nouhaila Innan and Muhammad Shafique

Title: Digital–analog quantum learning on Rydberg atom arrays

Abstract: We propose hybrid digital–analog (DA) learning algorithms on Rydberg atom arrays, combining the practical utility and near-term realizability of quantum learning with the scalable architectures of neutral atoms. Our approach requires only single-qubit operations in the digital setting and global driving according to the Rydberg Hamiltonian in the analog setting.
We perform a comprehensive numerical study of our algorithm on both classical and quantum data, using handwritten digit classification and unsupervised quantum phase boundary learning as representative tasks. In both cases, DA learning proves feasible in the near term, requires shorter circuit depths, and is more robust to realistic error models compared to fully digital learning schemes.
Our results indicate that DA learning provides a promising pathway for improved variational quantum learning experiments in the near term.
Authors: Jonathan Lu, Lucy Jiao, Kristina Wolinski, Milan Kornja?a, Hong-Ye Hu, Sergio Cantu, Fangli Liu, Susanne Yelin and Sheng-Tao Wang

Title: Distilling the knowledge with quantum neural networks

Abstract: Quantum Neural Networks (QNNs) are a promising class of quantum machine learning models with potential advantages when implemented on scalable, error-corrected quantum computers. However, as system sizes increase, deploying QNNs becomes challenging. Like their classical counterparts, a major obstacle is that large-scale QNNs may not be easily deployable on smaller systems with limited resources.
In this work, we address this challenge by compressing QNNs via knowledge distillation. We demonstrate how well-trained QNNs on large systems can be distilled into smaller architectures with similar configurations. Additionally, we explore the feasibility of transferring knowledge from classical neural networks to analogous QNNs.
Numerical results show that knowledge distillation reduces the training cost of QNNs in terms of qubit requirements and circuit depth. We also find that a self-knowledge-distillation approach can accelerate training convergence. These results provide new strategies for the efficient compression and practical deployment of QNNs.
Authors: Yuxuan Yan, Sitian Qian, Qi Zhao and Xingjian Zhang

Title: Double Descent in Quantum Kernel Methods

Abstract: The double descent phenomenon challenges traditional statistical learning theory by showing that larger models do not necessarily lead to reduced performance on unseen data. While this counterintuitive behavior has been observed in various classical machine learning models, particularly modern neural networks, it remains largely unexplored in quantum machine learning.
In this work, we analytically demonstrate that linear regression models in quantum feature spaces can exhibit double descent behavior, drawing on insights from classical linear regression and random matrix theory. Numerical experiments on quantum kernel methods across different

real-world datasets and system sizes further confirm the presence of a test error peak, a hallmark of double descent.
Our findings provide evidence that quantum models can operate in the overparameterized regime without overfitting, potentially enabling improved learning performance beyond the predictions of traditional statistical learning theory.
Authors: Marie Kempkes, Aroosa Ijaz, Elies Gil-Fuster, Carlos Bravo-Prieto, Jakob Spiegelberg, Evert van Nieuwenburg and Vedran Dunjko

Title: Double-bracket quantum algorithms for ground-state preparation via cooling

Abstract: Preparing ground states of Hamiltonians is a fundamental task in quantum computation with wide-ranging applications. While efficiently preparing approximate ground states on quantum hardware is challenging, nature achieves this naturally, inspiring thermodynamically motivated approaches such as imaginary-time evolution (ITE). However, synthesizing quantum circuits that efficiently implement such cooling methods remains difficult.
In this work, we propose cooling approaches for ground-state preparation by exploiting recently established Double-Bracket Quantum Algorithms (DBQA). Our contributions are two-fold:
We introduce Double-Bracket Quantum Imaginary-Time Evolution (DB-QITE), a new algorithm that compiles quantum circuits for ITE without requiring measurements. We provide rigorous guarantees that DB-QITE systematically lowers the energy of a state and increases its fidelity with the ground state.
We develop a more general framework called Double-Bracket Quantum Signal Processing (DB-QSP), which realizes shorter-depth circuits for ground-state preparation and extends to broader tasks involving polynomial transformations of Hamiltonians. This approach enables deterministic polynomial transformations of the Hamiltonian without auxiliary qubits or post-selection. We demonstrate the potential of DB-QSP to implement low-degree polynomial transformations as a “warm start” for existing QSP techniques that rely on post-selection.
Our algorithms are expected to serve both as standalone ground-state preparation methods in the early fault-tolerant era and as complementary tools alongside more established or heuristic approaches to ground-state preparation.
Authors: Yudai Suzuki, Bi Hong Tiang, Jeongrak Son, René Zander, Raphael Seidel, Nelly H. Y. Ng, Zoë Holmes and Marek Gluza

Title: Dynamical Regimes and Memory Performance in Quantum Reservoirs: Insights from Random Matrix Theory

Abstract: Quantum Reservoir Computing (QRC) provides a powerful framework for leveraging quantum dynamics to process information and is well-suited for various machine learning tasks. In this work, we investigate how different dynamical regimes—integrable, mixed, and chaotic—affect reservoir performance. Using simple quantum maps, such as the kicked rotor and kicked top, as prototype reservoirs, we find that chaotic dynamics, associated with thermalization, significantly enhance convergence, meaning that outputs primarily reflect recent input history. In contrast, integrable or localized regimes, constrained by conserved quantities, exhibit poor scrambling and slow convergence. Interestingly, quantum resonances in the kicked rotor produce superdiffusive transport, enabling faster information propagation than in the chaotic regime; however, rapid transport does not necessarily improve convergence, as resonant dynamics lack the strong scrambling needed for effective fading memory.
Performance is evaluated via learning tasks, with memory capacity quantified through the Information Processing Capacity (IPC). We identify the onset of thermalization—the crossover from

localized to ergodic behavior—as an optimal operating regime, where linear and nonlinear memory are well balanced, echoing the "edge of chaos" observed in classical reservoir computing. To benchmark performance across dynamical regimes, we employ random matrix transition ensembles that model interacting Floquet systems with level statistics interpolating between Poissonian and Wigner-Dyson distributions. This provides a universal upper bound on memory capacity, allowing us to pinpoint where IPC is maximized.
The kicked rotor and kicked top offer experimentally accessible, tunable platforms, where parameters such as kicking and coupling strengths provide precise control over dynamical complexity. These features make them promising candidates for implementation in cold atom systems and NMR-based quantum simulators.
Authors: Bharathi Kannan Jeevanandam, Nisarg Vyas, Sreeram Pg and Santhanam M.S

Title: EHands: Quantum Protocol for Polynomial Computation on Real-Valued Encoded States

Abstract: The novel constructive EHands protocol defines a universal set of reversible quantum operations for multivariable polynomial transformations on quantum processors. It introduces four basic subcircuits—multiplication, addition, negation, and parity flip—and uses expectation-value encoding (EVEN) to represent real numbers in quantum states. These elementary arithmetic operations can be systematically composed to compute degree-d polynomials, $P_d(x)$, on a quantum processing unit (QPU).
The resulting quantum circuit structure closely mirrors the stepwise evaluation of polynomials on a classical calculator, providing an intuitive and efficient approach to polynomial computation on quantum hardware. By enabling direct and predictable polynomial and nonlinear data transformations on a QPU, the method reduces dependence on classical post-processing in hybrid quantum-classical algorithms, facilitating advancements in many quantum applications.
The EHands quantum circuits are compact enough to deliver meaningful and accurate results on today's noisy quantum processors. We present a detailed implementation of $P_4(x)$ and report experimental results for polynomial approximations of common functions, obtained using IBM's Heron-class quantum processors and an ideal Qiskit simulator.
Authors: Jan Balewski, Chris Pestano, Mercy G. Amankwah, E. Wes Bethel and Talita Perciano

Title: Entanglement detection via machine learning techniques

Abstract: We present a deep learning framework for detecting quantum entanglement in bipartite systems, including bound entangled states. Our model combines a Convolutional Neural Network (CNN) with a Multilayer Perceptron (MLP) classifier within a one-class classification paradigm.
The training dataset includes bipartite separable states, negative partial transpose (NPT) entangled states, and synthetic noise samples drawn from a Gaussian distribution. Separable and NPT states are generated via random sampling according to the Hilbert-Schmidt measure, with NPT states selected using the Peres-Horodecki criterion.
Our model achieves a classification accuracy of 99.6% and demonstrates generalization capability by successfully identifying bound entangled states not seen during training. These results highlight the potential of artificial intelligence to address open challenges in quantum information theory and suggest new avenues for entanglement characterization beyond standard analytical criteria..
Authors: Daniel Uzcátegui, Katherine Muñoz, Aldo Delgado and Dardo Goyeneche

Title: Entanglement scaling in matrix product state representation of smooth functions and their shallow quantum circuit approximations

Abstract: Encoding classical data into a quantum state is a key component of many quantum algorithms. Recently, matrix product states (MPS) have emerged as a promising approach for constructing linearly deep quantum circuits that approximate input functions or distributions.
We derive rigorous, asymptotically tight bounds for entanglement decay in the MPS representation, depending on the smoothness of the complex input function. We show how this decay depends on localization and the function's support. Based on these insights, we construct an improved MPS-based algorithm that yields shallow and accurate quantum circuits for data encoding. Using Tensor Cross Interpolation, we build utility-scale circuits in a compute- and memory-efficient manner.
We validate our methods on heavy-tailed distributions relevant to finance and test the performance of our quantum circuits by executing and sampling from them on IBM quantum devices, for systems of up to 64 qubits.

Authors: Vladyslav Bohun, Illia Lukin, Mykola Lukhanko, Georgios Korpas, Philippe J.S. De Brouwer, Mykola Maksymenko and Maciej Koch-Janusz

Title: Entanglement-induced provable and robust quantum learning advantages

Abstract: Quantum computing has the potential to significantly enhance machine learning, yet a clear demonstration of quantum learning advantage has remained elusive. In this work, we rigorously establish a noise-robust, unconditional quantum learning advantage in expressivity, inference speed, and training efficiency compared to commonly used classical models.
Our proof is information-theoretic and identifies the source of this advantage: entanglement reduces the communication required for non-local tasks. Specifically, we design a task that can be solved with unit accuracy by quantum models using entanglement with a constant number of parameters, whereas commonly used classical models must scale linearly to achieve more than an exponentially small accuracy. The quantum model is also trainable with constant resources and remains robust against constant noise.
Through numerical simulations and trapped-ion experiments on IonQ Aria, we demonstrate this advantage. Our results provide guidance for realizing quantum learning advantages on current noisy intermediate-scale quantum (NISQ) devices.

Authors: Haimeng Zhao and Dong-Ling Deng

Title: Experimental quantum memristor-based reservoir computing

Abstract: Machine learning models have shown remarkable success across a wide range of problems, but their computational requirements can grow rapidly with task complexity, sometimes making them infeasible. To address this, two complementary approaches have emerged: biologically inspired neuromorphic computing, which seeks efficient learning processes, and quantum processing, which explores new computational paradigms.
A key challenge in combining these fields is achieving nonlinear responses, as quantum systems inherently evolve linearly. In this work, we tackle this challenge by designing and implementing a neuromorphic architecture on a novel photonic device: the quantum memristor. We perform nonlinear time series predictions and benchmark our model on four tasks, highlighting the essential role of the quantum memristive element. Our results demonstrate that it can serve as a building block for more sophisticated quantum neuromorphic networks.

Authors: Iris Agresti, Mirela Selimovic, Michal Siemaszko, Joshua Morris, Borivoje Dakic, Riccardo Albiero, Andrea Crespi, Francesco Ceccarelli, Roberto Osellame, Magdalena Stobinska and Philip Walther

Title: Exploring Trainability of Quantum Fourier Models for Different Data Re-uploading Schemes

Abstract: Quantum Machine Learning (QML) models show great promise for solving complex problems, but their expressivity is often constrained by the chosen architecture. Within the Fourier framework, layer-dependent data re-uploading has been proposed to enhance expressivity, yet this approach introduces a trade-off between generalizability and trainability. A large spectrum size can increase expressivity but may hinder training due to small frequency redundancies, whereas a small spectrum size limits generalizability.

In this work, we empirically investigate how data-tailored re-uploading schemes affect the trainability of QML models and identify indications for configurations that improve overall performance.

Authors: Felix Paul, Björn Minneker and Peter Jung

Title: Expressive equivalence of classical and quantum restricted Boltzmann machines

Abstract: Quantum computers hold the potential to efficiently sample from complex probability distributions, sparking growing interest in generative modeling within quantum machine learning. This has led to the development of numerous generative quantum models, though their trainability and scalability remain challenging. A notable example is the quantum restricted Boltzmann machine (QRBM), based on the Gibbs state of a parameterized non-commuting Hamiltonian. While expressive, QRBMs are computationally demanding for gradient evaluation, even on fault-tolerant quantum hardware.

In this work, we propose a semi-quantum restricted Boltzmann machine (sqRBM), designed for classical data to address these challenges. The sqRBM Hamiltonian is commuting in the visible subspace but non-commuting in the hidden subspace, enabling closed-form expressions for both output probabilities and gradients. This analytical tractability reveals a close relationship between sqRBMs and classical restricted Boltzmann machines (RBMs). Our theoretical analysis predicts that, to learn a given probability distribution, an RBM requires three times as many hidden units as an sqRBM, while both models maintain the same total number of parameters.

We validate these findings through numerical simulations involving up to 100 units. Our results suggest that sqRBMs could enable practical quantum machine learning applications in the near future by substantially reducing quantum resource requirements.

Authors: Maria Demidik, Cenk Tüysüz, Nico Piatkowski, Michele Grossi and Karl Jansen

Title: Expressivity Limits of Quantum Reservoir Computing

Abstract: Using non-digital physical systems as computational resources within the reservoir computing framework is motivated by two factors: the high-dimensional state space and the intrinsic nonlinear dynamics present in many physical substrates. Quantum systems, in particular, promise further advantages by offering an exponentially scaling state space relative to system size. However, how this exponential scaling translates into actual performance in quantum reservoir computers remains unclear.

Using methods from parameterized quantum circuits, we show that the expressive power of a quantum reservoir computer is not determined by the reservoir itself but is limited by the way input is injected. For the commonly used input encoding via single-qubit rotations, we provide an

upper bound on the maximum number of orthogonal functions the system can express. Notably, we find that expressivity scales linearly with the number of gates into which the input information is encoded, indicating that the exponentially large Hilbert space is not fully exploited for computation.
Authors: Nils-Erik Schütte, Niclas Götting, Hauke Müntinga, Meike List, Daniel Brunner and Christopher Gies

Title: Fake News Detection using Hybrid Classical-Quantum Transfer Learning Approach

Abstract: Introduction

Social media has transformed into a primary news source for younger audiences, often delivering information through catchy headlines and appealing visuals. However, users frequently consume this content without verifying its authenticity, making these platforms fertile ground for misinformation. Traditional fake news detection models primarily focus on textual features and may not generalize well across domains. Multimodal approaches have been proposed, but they typically rely on resource-heavy classical models. In response, we explore quantum-enhanced transfer learning to build a more generalizable and compact model.

Objectives

This work introduces QTL-FND, a hybrid classical-quantum model aimed at:

• Leveraging both textual and visual modalities for fake news detection,
• Exploiting the representational power of Variational Quantum Circuits (VQCs),
• Enhancing performance on low-resource, multi-domain datasets.

Methodology

Our model, QTL-FND, is a hybrid classical-quantum model that combines multimodal feature extraction with a quantum classifier. Text data is processed using a pre-trained DeBERTa model, while image features are extracted using ConvNeXt. The resulting embeddings are concatenated and passed through dense layers for dimensionality reduction. This fused vector is then fed into a 4-qubit dressed Variational Quantum Circuit (VQC) comprising parameterized rotation and entangling gates. The quantum output is measured and mapped to a binary class prediction. The model is trained and evaluated on a subset of the Fakeddit dataset containing 3823 real and 2919 fake news posts.

Results

On a test set of 1349 samples, QTL-FND achieves a test accuracy of 77.24%. The model shows good performance despite being trained on a smaller dataset with limited resources. Overall, the ability of the model to classify the news as real or fake is good, as indicated by the high number of true positives (343) and true negatives (699). However, when compared to a fully classical variant—where the quantum circuit was replaced with a dense layer—the classical model achieved slightly better accuracy (77.84%) and was more than 7× faster to train due to the overhead of simulating quantum circuits.

Discussion and Conclusion

The hybrid model does not outperform its classical counterpart, which showed slightly better accuracy and significantly faster execution. This aligns with prior findings12 where quantum models often lag behind in convergence and predictive power. The overhead of quantum circuit simulation remains a
bottleneck in practical use. Nonetheless, this work contributes to ongoing efforts to understand the potential and current limitations of quantum transfer learning. As hardware
and algorithms improve, hybrid quantum approaches may become more competitive in complex, data-limited domains.

Authors: Mitali Nanda and Tirupati Bolisetti

Title: Fast, Accurate and Interpretable Graph Classification with Topological Kernels: An Even More Scalable Alternative to Weisfeiler-Lehman Kernels

Abstract: We introduce a novel class of explicit feature maps that represent each graph by a compact feature vector, with entries corresponding to values of topological graph indices. These feature vectors are combined with a radial basis function (RBF) kernel to define a similarity measure between graphs.

We evaluate our approach on standard molecular datasets—PROTEINS, MUTAG, AIDS, DD, NCI1, and PTC-MR—and find that classification accuracies based on single topological-index feature vectors are lower than those of state-of-the-art substructure-based kernels. However, we observe a substantial reduction in kernel matrix computation time, up to an order of magnitude.

To improve representational capacity, we propose two extensions:

Extended Feature Vector (EFV): concatenating multiple topological indices into a single feature vector.

Linear Combination of Topological Kernels (LCTK): linearly combining RBF kernels computed on feature vectors of individual topological indices.

Empirical results show that these extensions yield additional accuracy gains of 5–12% across all datasets. End-to-end runtime comparisons against the Weisfeiler–Lehman subtree kernel reveal up to 15× speedup in Gram matrix construction and up to 3× speedup for feature vector construction on large random graphs.

We provide a detailed complexity analysis under both classical and quantum computational models, highlighting that computation of the Estrada index is amenable to exponential quantum speedup. Overall, our results demonstrate that low-dimensional representations based on topological graph indices provide both classical and quantum computational efficiency while maintaining high classification accuracy. These methods break the typical trade-off between computational time and predictive performance.

Authors: Adam Wesolowski, Ronin Wu and Karim Essafi

Title: Flexible quantum Kolmogorov-Arnold networks via generalized fractional-order Chebyshev functions and trainable QSVT basis

Abstract: Quantum neural networks (QNNs) have gained attention in recent years for exploring the practical utility of quantum computing. A recent addition to neural network architectures is the Kolmogorov-Arnold Network (KAN). Its key novelty over traditional neural networks is the enhanced freedom in choosing activation functions, allowing universality with model complexity scaling as O(m2)O(m^2)O(m2), where mmm is the input dimension, based on the Kolmogorov-Arnold theorem, which resolves Hilbert's thirteenth problem. In contrast, traditional neural networks achieve universality only in the limit of infinite size.

The quantum analogue of KANs (QKAN) encodes data as Block Encoded (BE) vectors, either from classical data or quantum states. QKAN transforms these inputs via learnable linear combinations of polynomial basis functions implemented through Quantum Singular Value Transformation (QSVT), enhancing the flexibility of activation functions. In the original QKAN, Chebyshev polynomials of the first kind are used as basis functions, yielding a model called CHEB-QKAN. The BE output can either be applied to a quantum state to extract classical outputs via Hadamard Tests or serve as state preparation. Other QKAN variants include VQKAN, Adaptive VQKAN (AVQKAN), and EVQKAN.

In this work, we focus on classical-to-classical QKAN models and make two main contributions:

Generalized Fractional-order Chebyshev Functions (GFCFs): a generalization of CHEB-QKAN to enhance adaptability.
Flexible QKAN (Flex-QKAN): a new QKAN variant with learnable basis functions via trainable QSVT angles, further improving flexibility.
Using CHEB-QKAN as the baseline, we compare our methods against other state-of-the-art QKANs.
Open-source implementations of all proposed variants, including combinations with GFCF and Flex-QKAN, are available at GitHub.
Authors: Javier González Otero, Adrián Pérez Salinas and Miguel Ángel González Ballester

Title: Forecasting the Lorenz System Using a Hierarchical Tensor Network Model

Abstract: We present a tensor network-based forecasting model for chaotic dynamical systems, demonstrated on the Lorenz system. The model uses a hierarchical tensor network architecture to capture non-Markovian temporal dependencies through multi-linear contractions. Trained on numerically generated Lorenz trajectories, it achieves accurate short-term reconstruction and forecasting, while longer-term predictions naturally diverge due to the system's chaotic sensitivity.
We systematically analyze the impact of bond dimension and tensor parameter homogeneity, finding that larger bond dimensions and inhomogeneous parametrizations improve expressivity, convergence speed, and generalization performance. These results highlight the potential of tensor network models as expressive and efficient tools for learning and forecasting complex nonlinear dynamics.
Authors: Jiabin You, Jian Feng Kong and Jun Ye

Title: From Bits to Qubits: Comparative Insights into Embedding Strategies for Machine Learning

Abstract: Quantum embedding is a key component in Quantum Machine Learning (QML), enabling the transformation of classical data into quantum states for efficient learning. This work provides a concise review of major quantum encoding strategies, including basis, amplitude, angle, QSample, associative memory (QuAM), QRAM, superdense, and Hamiltonian evolution encoding, and highlights their roles in QML tasks.
We systematically compare these methods in terms of qubit efficiency, circuit depth, and runtime complexity, offering practical guidance for selecting appropriate strategies based on dataset structure and hardware constraints. Special attention is given to the trade-offs between expressivity and implementability on Noisy Intermediate-Scale Quantum (NISQ) devices.
Building on insights from our recent IEEE Access publication (DOI: 10.1109/ACCESS.2024.3382150), we frame the discussion around QML-specific challenges, including hybrid encoding strategies, kernel-based learning, quantum data compression, and encoding-decoding efficiency. We also emphasize the increasing importance of encoding-aware design in quantum neural networks and quantum kernel methods.
Our contribution is twofold:
Provide a comparative synthesis of quantum embeddings from a QML perspective.
Serve as a roadmap for encoding strategy selection to optimize learning performance and resource utilization.
This work supports the development of scalable, hardware-adaptive QML pipelines.
Authors: Mansoor Ali Khan, Muhammad Naveed Aman and Biplab Sikdar

Title: Generalization Bounds for Quantum Learning via Rényi Divergences

Abstract: This study advances the theoretical understanding of quantum learning by introducing a new family of upper bounds on the expected generalization error of quantum learning algorithms. Building on the framework by Caro et al. (2024) and a new definition of the expected true loss, the work derives these bounds in terms of quantum and classical Rényi divergences.

A key innovation is the use of a variational technique to evaluate quantum Rényi divergences, including the Petz divergence and a newly proposed modified sandwich quantum Rényi divergence. Both analytical and numerical results show that bounds based on the modified sandwich divergence outperform those derived from the Petz divergence.

Additionally, the study provides probabilistic generalization error bounds via two approaches: one using the modified sandwich quantum Rényi divergence alongside classical Rényi divergence, and another based on the smooth max Rényi divergence. These results offer refined tools for assessing and improving the generalization performance of quantum learning algorithms.

Authors: Naqueeb Ahmad Warsi, Ayanava Dasgupta and Masahito Hayashi

Title: Generalization Bounds in Hybrid-Quantum Machine Learning Models

Abstract: Hybrid quantum-classical models aim to harness the strengths of both quantum computing and classical machine learning, but their practical potential and generalization capabilities remain poorly understood. A comprehensive theoretical understanding of their ability to generalize

from finite training data is essential for guiding their development and identifying scenarios for genuine quantum advantage.

This work addresses this gap by developing a unified mathematical framework for analyzing generalization in such hybrid models. We establish a novel generalization bound for a hybrid model trained on N data points, with a quantum part comprising T trainable quantum gates and a classical part consisting of k bounded fully-connected layers. The generalization error is bounded with high probability by O-tilde(sqrt((T*log(T))/N) + (alpha^k)/sqrt(N)), where the norms of the classical layers are bounded by alpha.

This bound decomposes cleanly into distinct quantum and classical contributions, extending prior work and clarifying their interaction. Our analysis indicates that for a small number of classical layers, generalization is predominantly influenced by the quantum component's complexity. However, as the number of classical layers increases, the classical term can dominate the generalization bound. Our derivation uses covering numbers to quantify the complexity of the quantum and classical components separately, then combines these measures using Dudley's entropy integral to bound the Rademacher complexity.

This result shows that introducing bounded classical layers on top of a trainable quantum model does not necessarily degrade generalization performance. Instead, hybrid architectures can retain the learning guarantees of their fully quantum counterparts while offering practical benefits, such as reduced quantum circuit depth. This theoretical insight is a crucial step towards the principled design, evaluation, and deployment of hybrid quantum algorithms in practical machine learning scenarios.

Authors: Tongyan Wu, Amine Bentellis, Alona Sakhnenko and Jeanette Lorenz

Title: Gradient Scalability on Super-polynomially Complex Quantum Landscapes

Abstract: Variational Quantum Algorithms (VQAs) are promising for near-term quantum computing, but they face scalability challenges due to barren plateaus, where gradients vanish

exponentially with system size. Some conjectures suggest that avoiding barren plateaus may inherently make VQAs classically simulable, potentially limiting quantum advantage.
This work advances the theoretical understanding of the link between trainability and computational complexity in VQAs:
Introduces the Linear Clifford Encoder (LCE), a technique that maintains constant-scaling gradient statistics in regions near Clifford circuits.
Uses classical Taylor surrogates to reveal computational complexity phase transitions, from polynomial to super-polynomial, as the size of the initialization region increases.
Numerical experiments on LCE-transformed landscapes suggest a super-polynomially complex transition region where gradients decay only polynomially.
These results highlight a connection between trainability and computational complexity, suggesting a plausible path to barren-plateau-free quantum advantage by carefully choosing initialization regions and leveraging LCE transformations.
Authors: Sabri Meyer, Francesco Scala, Franceso Tacchino and Aurelien Lucchi

Title: Grover's algorithm with W state-based initialization for solving the exact-cover problem

Abstract: Grover's algorithm provides a quadratic speedup over classical brute-force search in terms of query complexity for unstructured search problems, and its potential applications are actively being explored. In this work, we focus on the exact-cover problem, which is a well-known NP-complete problem, and explore how Grover's algorithm can be applied to solve it more efficiently using problem-specific initial states. While Grover's algorithm typically begins with a uniform superposition state, we
propose an alternative initialization using a W state that is tailored to the exact-cover problem. We show that this customized initialization remarkably reduces the number of queries, and validate its effectiveness by comparing quantum resource requirements in an existing Grover-based algorithm for the exact-cover problem. These results highlight the importance of problem-specific initialization in improving the resource efficiency and practicality of Grover-based algorithms.
Authors: Eunok Bae, Nari Choi, Jeonghyun Shin and Minjin Choi

Title: Hardware Adapted Quantum Machine Learning with Pulse-Level optimization

Abstract: Quantum Machine Learning often overlooks hardware constraints, relying on idealized gate-based abstractions. We introduce Pulsed Quantum Machine Learning, a framework that replaces parameterized gates with native quantum pulses to align more closely with physical hardware. As a case study, we adapt a data re-uploading model by introducing a pulse-based encoding that preserves dataset geometry and enables direct, hardware-native control. Simulations on both synthetic and real datasets demonstrate improved performance and robustness over gate-based models, particularly under noise and increasing circuit depth.
Authors: Ignacio Acedo

Title: Harnessing quantum back-action for time-series processing

Abstract: Quantum measurements fundamentally affect quantum systems through back-action, a phenomenon that has been extensively studied in quantum mechanics. While projective measurements extract maximal classical information, they significantly alter the system state, potentially disrupting the quantum computation process. In contrast, weak measurements offer a delicate balance between information extraction and system disturbance, presenting an intriguing avenue for quantum information processing. In this work, we demonstrate that incorporating weak

measurements into a quantum machine learning protocol known as quantum reservoir computing provides advantages in both execution time scaling and overall performance.
Authors: Giacomo Franceschetto, Marcin P?odzie?, Maciej Lewenstein, Antonio Acín and Pere Mujal

Title: Heuristic ansatz design for trainable ion-native digital-analog quantum circuits
Abstract: Variational quantum algorithms have become a standard tool for modern-day quantum computing. When designing an appropriate ansatz configuration for specific problems, a possible approach to account for hardware specifications comes in the form of digital-analog quantum circuits, where sequences of quantum gates are alternated with entangling natural Hamiltonian evolution. We consider these challenges for the example of ion-based quantum computers, where a hardware-native circuit has recently been proposed for the Quantum Approximate Optimization Algorithm. In our work, we propose a heuristic for identifying parameters of the natural Hamiltonian evolution that can significantly boost circuit trainability. This approach allows notably reducing the number of required circuit layers, bringing the algorithm one step closer to practical implementation.
Authors: Daniil Rabinovich, Luis Ernesto Campos Espinoza, Georgii Paradezhenko and Kirill Lakhmanskiy

Title: Hybrid Parameterized Quantum States for Variational Quantum Learning
Abstract: Variational quantum learning faces key challenges in the noisy intermediate-scale quantum era. Parameterized quantum circuit models suffer from finite-shot uncertainty and quantum noise, while neural quantum states lack genuine quantum correlations and scalability. We propose Hybrid Parameterized Quantum States, a flexible framework that blends PQC-based measurements with neural estimators via postprocessing functions for shot-efficient learning under hardware constraints. HPQS is demonstrated on three tasks: (1) Expectation-based QML, achieving higher accuracy than PQC-only and NQS-only baselines; (2) Quantum-Train, generating classical network parameters with polylogarithmic variables; and (3) Quantum Parameter Adaptation, producing LoRA adapters for fine-tuning large language models (GPT-2, Gemma-2) under low-shot conditions. These results position HPQS as a scalable, noise-resilient framework for variational quantum learning on current and future quantum hardware.
Authors: Chen-Yu Liu

Title: Hybrid Quantum Kolmogorov-Arnold Networks for High Energy Physics Analysis at the LHC
Abstract: The advent of the High-Luminosity Large Hadron Collider presents significant computational challenges, requiring innovative machine learning architectures for High Energy Physics data analysis. Kolmogorov-Arnold Networks, inspired by the Kolmogorov-Arnold representation theorem, differ from traditional Multi-Layer Perceptrons by using learnable activation functions on network edges instead of fixed activations on nodes. This work introduces Hybrid Quantum KANs, a hybrid quantum-classical architecture where the base linear transformation component of a KAN layer is implemented by a variational quantum circuit. We describe the architectural differences between KANs and QKANs and present a comparative performance analysis on a benchmark HEP task: the classification of high-pT jets using simulated LHC proton-proton collision data via the Pennylane framework. Preliminary results indicate that QKANs achieve validation performance comparable to classical KANs, suggesting potential for improved generalization and further exploration in complex scientific applications.

Code available at: https://github.com/elucidator8918/QKAN-ML4SCI/tree/main/Task-IX
Authors: Siddhant Dutta and Sadok Ben Yahia

Title: Hybrid Quantum Transfer Learning Models for Credit Risk Assessment

Abstract: In this work, we explore the potential of hybrid quantum-classical transfer learning models in enhancing credit risk classification performance. We analyze two quantum transfer learning strategies that integrate a classical neural network with a parameterized quantum circuit, leveraging pretraining and co-training techniques respectively. The first model utilizes a frozen classical feature extractor with a quantum refinement layer trained on task-specific financial data. The second model co-trains both the classical and quantum components end-to-end, enabling synergistic adaptation across the architecture. Experimental evaluations on a real-world normalized credit risk dataset demonstrate that both models achieve improved classification metrics over a purely classical baseline. Notably, the co-trained hybrid model achieves a higher area under the ROC curve and improved recall for the minority class, highlighting its capacity to address class imbalance challenges. While the co-trained model exhibits increased computational overhead, the results substantiate the viability of quantum transfer learning approaches for real-world financial decision-making under resource-aware conditions.
Authors: Parvathy Gopakumar, Rubell Marion Lincy G, Salvatore Sinno and Shruthi Thuravakkath

Title: Hybrid quantum-classical heuristics for optimizing large separable operators

Abstract: We study heuristic algorithms for linear optimization problems over the cone of separable operators. This is a challenging problem since determining separability is in general NP-hard and the dimension grows exponentially with the number of qubits. We offer a solution to both of these problems, one by introducing a heuristic algorithm that dramatically reduces the dimension of the problem by using a quantum co-processor. We also show numerically that see-saw algorithms perform well when the dimension of the problem is not too large. An important feature of these algorithms is that they yield feasible solutions, not just bounds on the optimal value. We apply our algorithm to the Hamiltonian problem, where we wish to find the separable states of least energy. By comparing this to the ground energy, we are able to define a measure of entanglement for the groundspace of a given Hamiltonian.
Authors: Ankith Mohan, Tobias Haug, Kishor Bharti and Jamie Sikora

Title: Hybrid Quantum-Classical Traffic Flow Classification with Deep Feature Extraction

Abstract: We propose a hybrid quantum–classical model for multi-class urban traffic flow classification under real-world constraints. The model, termed Deep Neural Network–Variational Quantum Classifier, combines a classical deep neural network for feature extraction with variational quantum circuits for multi-class classification. Four traffic attributes—speed, volume, occupancy, and vehicle size distribution—are encoded into quantum states using amplitude or angle encoding schemes.
By leveraging discriminative features extracted by the neural network, the model improves the effectiveness and robustness of quantum classification compared to standalone variational quantum circuits. All experiments are conducted in simulation using Qiskit's StatevectorSampler on an urban Vehicle Detection System dataset. The model achieves a test accuracy of 80.0%, outperforming the VQC-only baseline of 63.9% while remaining compatible with current noisy intermediate-scale quantum devices.

These results highlight the benefits of classical preprocessing in quantum learning pipelines and motivate further research on hybrid architectures, efficient encoding strategies, and noise-resilient circuit designs for near-term quantum applications.
Authors: Hongsuk Yi

Title: Implementing Quantum Transformers on Trapped Ion Devices
Abstract: The transformer architecture has become a cornerstone of modern sequential machine learning, revolutionizing fields from natural language processing and computer vision to computational genomics. Its success is largely attributed to the self-attention mechanism, a key subroutine that captures long-range correlations within input sequences. However, the computation of self-attention uses extensive computational resources. As classical hardware approaches its physical limits, the rapid progress in quantum computing presents a new paradigm for overcoming such computational challenges. In this work, we focus on the practical implementation and evaluation of two leading quantum transformer models on state-of-the-art quantum hardware. Our study provides an account of the challenges and opportunities encountered when deploying these models on noisy near-term quantum devices, highlighting both the practical feasibility and current limitations of quantum transformers.
Authors: Zhan Yu, Naixu Guo, Gabriel Matos, Nikhil Khatri, Pranav Kalidindi, Yizhan Han, Lirandë Pira, Soumik Adhikary, Steve Clark and Patrick Rebentrost

Title: Information plane and compression-gnostic feedback in quantum machine learning
Abstract: The information plane has been proposed as an analytical tool for studying the learning dynamics of neural networks. It provides quantitative insight on how the model approaches the learned state by approximating a minimal sufficient statistics. In this paper, we extend this tool to the domain of quantum learning models. In a second step, we study how the insight on how much the model compresses the input data can be used to improve a learning algorithm. Specifically, we consider two ways to do so: via a multiplicative regularization of the loss function, or with a compression-agnostic scheduler of the learning rate for algorithms based on gradient descent. Both ways turn out to be equivalent in our implementation. Finally, we benchmark the proposed learning algorithms on several classification and regression tasks using variational quantum circuits. The results demonstrate an improvement in test accuracy and convergence speed for both synthetic and real-world datasets. Additionally, with one example we analyzed the impact of the proposed modifications on the performance of neural networks in a classification task.
Authors: Nathan Haboury, Mo Kordzanganeh, Alexey Melnikov and Pavel Sekatski

Title: Interpretable Machine Learning for Quantum Control
Abstract: Interpretable machine learning is a new direction aiming at understanding the behaviour of blackbox complex structures such as neural networks. Recently, there have been efforts to establish such frameworks in quantum machine learning. In the broader machine-learning community, however, recent studies have emphasized the distinction between intrinsically interpretable models and post-hoc explainable approaches. In this work, we present a hybrid interpretable–explainable framework for quantum control applications. The objective is to extract structured insight into how control pulses influence quantum system dynamics in the presence of noise. At its core lies a local analytic expansion that quantifies the system’s response to small pulse perturbations, revealing how deviations from ideal behaviour arise from environmental noise and control imperfections. Our approach builds on previous single-qubit and qudit graybox

applications, which demonstrated high-fidelity gate synthesis. We applied a new framework to a qutrit system under strong noise, showing why our optimized pulses exhibit dramatically reduced noise sensitivity. This bridges the gap between abstract model outputs and actionable insights, laying a practical foundation for interpretable quantum control in noisy, high-dimensional settings.
Authors: Yule Mayevsky, Akram Youssry, Ritik Sareen, Gerardo Paz-Silva and Alberto Peruzzo

Title: Is data-efficient learning feasible with quantum models?

Abstract: The importance of analyzing nontrivial datasets when testing quantum machine learning models is becoming increasingly prominent in literature, yet a cohesive framework for understanding dataset characteristics remains elusive. In this work, we concentrate on the size of the dataset as an indicator of its complexity and explore the potential for quantum machine learning models to demonstrate superior data efficiency compared to classical models, particularly through the lens of quantum kernel methods. We provide proof of the existence of classical datasets where quantum kernel methods achieve data efficiency by generating semi-artificial datasets. Additionally, our study introduces a new analytical tool to the quantum machine learning domain, derived from classical kernel methods, which can be utilized for investigating the classical-quantum gap. Our results pave a way to a comprehensive exploration of dataset complexities, providing insights into how these complexities influence quantum machine learning performance relative to traditional methods. This research contributes to a deeper understanding of the generalization benefits of quantum machine learning models, setting the stage for future advancements in the field.
Authors: Alona Sakhnenko, Christian Mendl and Jeanette Miriam Lorenz

Title: Learning Boolean Functions with Non-Local Dependencies via Hybrid Quantum-Classical Neural Networks

Abstract: Quantum Machine Learning presents a compelling avenue for addressing complex computational challenges by integrating quantum resources into classical machine learning paradigms. This work details a hybrid quantum-classical neural network designed for binary classification, specifically applied to the Subset Weight Comparison Problem. This novel task challenges the network to discern a binary output based on the relative Hamming weights within two distinct, disjoint subsets of an input binary string. Formally, for a binary string b in {0, 1}^N and two disjoint subsets of indices S1, S2 ⊂ {0, 1, …, N − 1}, the function C_S1,S2(b) is defined as 1 if the sum of b_j over j in S1 is greater than the sum of b_j over j in S2, and 0 otherwise.
For instance, with N = 10, S1 = {0, 2, 5, 8}, and S2 = {1, 3, 6, 9}, the input b = [0, 0, 1, 1, 0, 0, 0, 1, 0, 1] yields the sum over S1 as 0 + 1 + 0 + 0 = 1 and the sum over S2 as 0 + 1 + 0 + 1 = 2. Since 1 > 2 is false, C_S1,S2(b) = 0. This problem, unlike simpler parity or single-subset majority functions, necessitates the learning of intricate non-local dependencies and a dynamic comparative decision boundary, making it a more challenging benchmark for quantum machine learning architectures.
Authors: Abdullah Kazi and Jayesh Hire

Title: Learning to erase quantum states: thermodynamic implications of quantum learning theory

Abstract: The energy cost of erasing quantum states depends on our knowledge of the states. We show that learning algorithms can acquire such knowledge to erase many copies of an unknown state at the optimal energy cost. This is proved by showing that learning can be made fully reversible and has no fundamental energy cost itself. With simple counting arguments, we relate

the energy cost of erasing quantum states to their complexity, entanglement, and magic. We further show that the constructed erasure protocol is computationally efficient when learning is efficient. Conversely, under standard cryptographic assumptions, we prove that the optimal energy cost cannot be achieved efficiently in general. These results also enable efficient work extraction based on learning. Together, our results establish a concrete connection between quantum learning theory and thermodynamics, highlighting the physical significance of learning processes and enabling efficient learning-based protocols for thermodynamic tasks.
Authors: Haimeng Zhao, Yuzhen Zhang and John Preskill

Title: Learning to generate high-dimensional distributions with low-dimensional quantum Boltzmann machines

Abstract: In recent years, researchers have been exploring ways to generalize Boltzmann machines to quantum systems, leading to the development of variations such as fully-visible and restricted quantum Boltzmann machines. Due to the non-commuting nature of their Hamiltonians, restricted quantum Boltzmann machines face trainability issues, whereas fully-visible quantum Boltzmann machines have emerged as a more tractable option, as recent results demonstrate their sample-efficient trainability. These results position fully-visible quantum Boltzmann machines as a favorable choice, offering potential improvements over fully-visible Boltzmann machines without suffering from the trainability issues associated with restricted quantum Boltzmann machines. In this work, we show that low-dimensional, fully-visible quantum Boltzmann machines can learn to generate distributions typically associated with higher-dimensional systems. We validate our findings through numerical experiments on both artificial datasets and real-world examples from the high energy physics problem of jet event generation. We find that non-commuting terms and Hamiltonian connectivity improve the learning capabilities of quantum Boltzmann machines, providing flexible resources suitable for various hardware architectures. Furthermore, we provide strategies and future directions to maximize the learning capacity of fully-visible quantum Boltzmann machines.
Authors: Cenk Tüysüz, Maria Demidik, Luuk Coopmans, Enrico Rinaldi, Vincent Croft, Yacine Haddad, Matthias Rosenkranz and Karl Jansen

Title: Learning Unitaries with Quantum Statistical Queries

Abstract: We propose several algorithms for learning unitary operators from quantum statistical queries with respect to their Choi-Jamiolkowski state. Quantum statistical queries capture the capabilities of a learner with limited quantum resources, which receives as input only noisy estimates of expected values of measurements. Our approach leverages quantum statistical queries to estimate the Fourier mass of a unitary on a subset of Pauli strings, generalizing previous techniques developed for uniform quantum examples. Specifically, we show that the celebrated quantum Goldreich-Levin algorithm can be implemented with quantum statistical queries, whereas the prior version of the algorithm involves oracle access to the unitary and its inverse. As an application, we prove that quantum Boolean functions with constant total influence or with constant degree are efficiently learnable in our model. Moreover, we prove that log(n)-juntas are efficiently learnable and constant-depth circuits are learnable query-efficiently with quantum statistical queries. On the other hand, all previous algorithms for these tasks demand significantly greater resources, such as oracle access to the unitary or direct access to the Choi-Jamiolkowski state.
We also demonstrate that, despite these positive results, quantum statistical queries lead to an exponentially larger query complexity for certain tasks, compared to separable measurements of

the Choi-Jamiolkowski state. In particular, we show an exponential lower bound for learning a class of phase-oracle unitaries and a double exponential lower bound for testing the unitarity of channels. Taken together, our results indicate that quantum statistical queries offer a unified framework for various unitary learning tasks, with potential applications in quantum machine learning, many-body physics, and benchmarking of near-term devices.
Authors: Armando Angrisani

Title: Low Cost Experimental Design for Frequency Estimation

Abstract: Frequency estimation is a crucial task in quantum metrology. It is relevant for a wide range of physical phenomena, such as Larmor, Rabi, and Ramsey oscillations, with applications in sensing and calibration, as well as quantum problems, namely phase and amplitude estimation. Bayesian inference can be applied to this problem and has been shown capable of saturating the Heisenberg limit. However, the optimization costs are high. For this reason, lightweight heuristics have been widely adopted.

In this work, we thoroughly evaluate the performance of these heuristics for frequency estimation, comparing them with the fundamental limits of metrology and a reference random strategy. We explore the advantages and shortcomings of these approaches and propose two lower-cost and more stable adaptive algorithms. Our methods are based on WES, a window expansion strategy that drives an adaptive problem-tailored definition of the search range, along with other cost-cutting measures. They allow extra classical resources to be traded in for increased quantum advantage. One of the methods seeks variance minimization, while the other considers only a measure of statistical efficiency, making it more general and robust, for example in view of multi-modality.

We benchmark our algorithms against all others in ideal and noisy scenarios, showing that they achieve the most reliable performance and fastest learning rate, saturating the Heisenberg limit.
Authors: Alexandra Ramôa, Luis Santos and Akihito Soeda

Title: Machine Learning-Assisted Parametric Modulation in Atomic Magnetometry

Abstract: Atomic magnetometers have garnered significant attention due to their exceptional sensitivity and broad applicability in physics, biomedical diagnostics, and geophysics. Among these, vector magnetometers, capable of measuring magnetic fields along multiple axes, are particularly valuable. We introduce a novel machine learning-assisted approach to parametric modulation-based vector atomic magnetometry, aimed at overcoming critical limitations of previous methodologies, including high crosstalk between orthogonal components and limited measurement of magnetic field components to only two orthogonal axes. Utilizing deep neural networks, our method substantially minimizes inter-axis crosstalk, potentially enhances the sensitivity of our sensor, and expands measurement capabilities to three-dimensional vector magnetic field detection. This advancement sets a new performance benchmark for vector magnetometry and demonstrates the powerful potential of integrating artificial intelligence to optimize and simplify complex sensor systems.
Authors: Vincent Han Leong Lau, Wenyu Guo, Ratnajit Sarkar, Aaron Tranter, Qian Ling Kee, Rigui Zhou, Ping Koy Lam, Mile Gu and Tao Wang

Title: MANTIS: Multiple Anomaly-Detection Networks for Tensor Inspired Solutions

Abstract: Anomaly detection is vital across a wide range of applications, from cybersecurity to quality assurance, where the goal is to identify fraudulent activities or unexpected patterns in data.

A central challenge in this task is that models are typically trained only on normal data, while anomalies—by nature—are diverse and sparse, occupying a virtually unbounded space. With advanced technologies, neural networks have been used for detection and classification, but explainability remains a challenge. Tensor networks, originally developed in quantum many-body physics, have been identified recently as an alternative framework, providing additional benefits of efficiency and scalability. By leveraging their ability to compactly represent high-order correlations in data, tensor networks can serve as interpretable and data-efficient models.
In this work, we train a new model of tensor networks defined as the superposition of multiple bond dimension-one matrix product operators. The novel model that we implement is advantageous as it is highly parallelizable, explainable, and lightweight. We demonstrate how tensor network representations can be trained to model the typical behavior of a dataset and subsequently identify anomalies with features that deviate significantly from the learned structure. Our approach shows promising performance on benchmark datasets, achieving competitive accuracy while offering significant advantages in terms of model size and interpretability. These results suggest that superpositions of bond-dimension one tensor networks offer a new physics-inspired toolkit for scalable and explainable anomaly detection in complex data environments.
Authors: Si Min Chan, Apimuk Sornsaeng and Dario Poletti

Title: More-efficient Quantum Multivariate Mean Value Estimator from Generalized Grover Operator

Abstract: In this work, we present an efficient algorithm for multivariate mean value estimation. Our algorithm outperforms previous work by polylog factors and nearly saturates the known lower bound. More formally, given a random vector of dimension d, we find an algorithm that uses O(n log(d/δ)) samples to find a mean estimate that differs from the true mean by √(tr Σ)/n in ℓ∞ norm and hence √(d tr Σ)/n in ℓ2 norm, where Σ is the covariance matrix of the components of the random vector. We also present another algorithm that uses smaller memory but costs an extra d^(1/4) in complexity.
Consider the Grover operator, the unitary operator used in Grover's algorithm. It contains an oracle that uses a ±1 phase for each candidate in the search space. Previous work has demonstrated that when we substitute the oracle in the Grover operator with generic phases, it can serve as a good mean value estimator in some mathematical sense. We used this idea to build our algorithm. Our result is not exactly optimal due to a log(d/δ) term in our complexity, as opposed to something nicer such as log(1/δ); this comes from the phase estimation primitive in our algorithm. So far, this primitive is the only major known method to tackle the problem, and moving beyond this idea seems hard. Our results demonstrate that the methodology with generalized Grover operators can be used to develop the optimal algorithm without polylog overhead for different tasks relating to mean value estimation.
Authors: Letian Tang

Title: Multiple photons enhance data efficiency in quantum machine learning

Abstract: Machine learning has delivered transformative capabilities across science and technology, while a common bottleneck is the large amount of data required to train models. Recently, quantum machine learning has emerged as a promising approach to use the features of quantum mechanics to enhance machine learning. A powerful platform for quantum machine learning is photonic quantum information processing, which harnesses the intrinsic robustness and long coherence times of photons at room temperature.

Here, we show that multi-photon states propagating through linear-optical circuits learn from data more effectively, achieving a provable advantage in the number of training data required compared to single-photon or coherent states. For the tasks of unitary and metric learning, we demonstrate that multi-photon protocols reach higher test accuracy while requiring significantly smaller training datasets. We implement these protocols experimentally on a fully programmable photonic integrated platform and introduce a quantum geometric framework to rigorously characterize performance and resource scaling.
Our results open new directions for practical and data-efficient applications of photonic quantum machine learning.
Authors: Yong Wang, Zhenghao Yin, Tobias Haug, Ciro Pentangelo, Simone Piacentini, Andrea Crespi, Francesco Ceccarelli, Roberto Osellame and Philip Walther

Title: Near-Optimal Parameter Tuning of Level-1 QAOA for Ising Models

Abstract:The Quantum Approximate Optimisation Algorithm is a hybrid quantum-classical algorithm for solving combinatorial optimisation problems. QAOA encodes solutions into the ground state of a Hamiltonian, approximated by a p-level parameterised quantum circuit composed of problem and mixer Hamiltonians, with parameters optimised classically. While deeper QAOA circuits can offer greater accuracy, practical applications are constrained by complex parameter optimisation and physical limitations such as gate noise, restricted qubit connectivity, and state-preparation-and-measurement errors, limiting implementations to shallow depths.
This work focuses on QAOA at p=1 for QUBO problems, represented as Ising models. Despite having only two parameters, gamma and beta, we show that their optimisation is challenging due to a highly oscillatory landscape, with oscillation rates increasing with the problem size, density, and weight. This behaviour necessitates high-resolution grid searches to avoid distortion of cost landscapes that may result in inaccurate minima. We propose an efficient optimisation strategy that reduces the two-dimensional gamma-beta search to a one-dimensional search over gamma, with beta* computed analytically. We establish the maximum permissible sampling period required to accurately map the gamma landscape and provide an algorithm to estimate the optimal parameters in polynomial time.
Furthermore, we rigorously prove that for regular graphs on average, the globally optimal gamma* values are concentrated very close to zero and coincide with the first local optimum, enabling gradient descent to replace exhaustive line searches. This approach is validated using Recursive QAOA, where it consistently outperforms both coarsely optimised RQAOA and semidefinite programs across all tested QUBO instances.
Authors: V Vijendran, Dax Enshan Koh, Eunok Bae, Hyukjoon Kwon, Ping Koy Lam and Syed M Assad

Title: Neural Quantum Embedded Self-supervised Learning

Abstract: Self-supervised learning aims to learn meaningful data representations without requiring labeled data. SSL algorithms based on contrastive learning have achieved significant success across various domains. Recently, several approaches have integrated parameterized quantum circuits into SSL frameworks. In this work, inspired by Neural Quantum Embedding, we propose a quantum-classical hybrid method for representation learning, in which a neural network embeds data into quantum states, enabling the use of quantum state fidelity as a contrastive loss. Unlike conventional methods that measure similarity in the n-dimensional Euclidean space, our approach compares n-dimensional representation vectors in the $(2n - 1)$-dimensional Hilbert space using quantum state fidelity.

Specifically, we introduce two contrastive losses: (1) a quantum state fidelity-based variant of the NT-Xent loss, and (2) a squared quantum state fidelity loss with pseudo labels. We evaluate the learned representations on the CIFAR-10 dataset using SimCLR, QSSL, and our proposed models, all sharing the same CNN backbone. We assess performance using both a linear classifier (single-layer perceptron) and a quantum classifier (quantum convolutional neural network). Experimental results, validated by the Mann–Whitney U test, show that in the 8-qubit quantum embedding setting, our second proposed model significantly outperforms SimCLR in linear classification, and both SimCLR and QSSL in quantum classification. However, in the 4-qubit quantum embedding setting, our models do not show statistically significant improvement over the baselines. These results suggest that leveraging a sufficiently large quantum representation space can enhance the effectiveness of self-supervised learning.
Authors: Dohyoung Lee, Daniel K. Park and Taeyoung Park

Title: New aspects of quantum topological data analysis: Betti number estimation, and testing and tracking of homology and cohomology classes

Abstract: Topological Data Analysis has emerged as a robust framework for extracting global structural features, such as connected components, loops, and voids, from high-dimensional data. Central to this methodology are Betti numbers, which count the number of k-dimensional topological holes in a simplicial complex, offering a compact summary of the data's shape. While classical algorithms for computing Betti numbers and persistent homology have seen widespread adoption, their computational cost becomes prohibitive for large or high-dimensional datasets. This has led to growing interest in quantum algorithms as a means to accelerate topological computations.

A foundational result in this direction is the quantum algorithm by Lloyd, Garnerone, and Zanardi, which estimates Betti numbers using quantum phase estimation applied to combinatorial Laplacians. This work spurred extensive research at the intersection of quantum computing, many-body physics, and computational topology. Subsequent studies have expanded the scope of quantum topological data analysis, but have also revealed fundamental computational hardness results that constrain the possibility of exponential quantum speedups in general settings.

In this work, we investigate new algorithmic and structural avenues in quantum topological data analysis, focusing on both homological and cohomological invariants. Our contributions are twofold. First, we introduce a novel input model for Betti number and persistent Betti number estimation, which provides structured quantum access to simplicial complexes. Unlike prior models that rely on explicit Laplacian construction and kernel dimension estimation, our approach circumvents the need to compute combinatorial Laplacians altogether. Instead, we develop a homology-tracking technique based on localized queries and combinatorial primitives, allowing for more efficient quantum estimation protocols. We show that this method achieves significant speedups over existing classical and quantum algorithms, and can yield exponential improvements under certain natural input distributions.

Second, we initiate the study of homology property testing in the quantum setting, proposing a suite of problems that capture finer-grained topological features beyond raw Betti numbers. We define and analyze several quantum algorithms for testing whether a given simplicial complex satisfies certain homological or cohomological properties. This includes algorithms for testing the triviality of cohomology classes and distinguishing between non-isomorphic classes, tasks that are computationally demanding classically. We demonstrate that these problems admit efficient quantum algorithms, often exhibiting exponential advantages in query and time complexity relative to classical baselines.

Collectively, our results reveal new possibilities for exploiting quantum resources in topological

data analysis. By developing quantum algorithms that go beyond Laplacian-based methods and incorporating property testing into the quantum topological data analysis framework, we open up novel pathways for both theoretical understanding and practical application. These findings underscore the potential of quantum computing as a powerful tool in the analysis of complex data through the lens of algebraic topology, and point toward broader opportunities for demonstrating quantum advantage in data-centric domains.
Authors: Junseo Lee and Nhat A. Nghiem

Title: Nonstabilizerness Enhances Thrifty Shadow Estimation
Abstract: Shadow estimation is a powerful approach for estimating the expectation values of many observables. Thrifty shadow estimation is a simple variant that is proposed to reduce the experimental overhead by reusing random circuits repeatedly. Although this idea is simple, its performance is quite elusive. In this work, we show that thrifty shadow estimation is effective on average whenever the unitary ensemble forms a 2-design, in sharp contrast with previous expectations. In thrifty shadow estimation based on the Clifford group, the variance is inversely correlated with the degree of nonstabilizerness of the state and observable, which is a key resource in quantum information processing. For fidelity estimation and purity estimation, it decreases exponentially with the stabilizer 2-Rényi entropy of the target state, which endows the stabilizer 2-Rényi entropy with a clear operational meaning. In addition, we propose a simple circuit to enhance the efficiency, which requires only one layer of T gates and is particularly appealing in the NISQ era.
Authors: Datong Chen and Huangjun Zhu

Title: On the Generalization of Adversarially Trained Quantum Classifiers
Abstract: Quantum classifiers are vulnerable to adversarial attacks perturbing their input classical or quantum data. Adversarial training has emerged as a promising countermeasure, where quantum classifiers are trained using an attack-aware loss function. We establish novel bounds on the generalization error of adversarially trained quantum classifiers in terms of the perturbation strength of the adversary. The bound quantifies the excess generalization error incurred to ensure robustness to adversarial attacks, informing us about the number of training samples needed to ensure good generalization, while also yielding insights into the impact of the quantum embedding. For quantum binary classifiers employing angle embedding, we find that, in the presence of adversarial attacks on classical inputs, the increase in sample complexity due to adversarial training over conventional training vanishes in the limit of high-dimensional inputs. We validate our theoretical findings with numerical experiments.
Authors: Petros Georgiou, Aaron Mark Thomas, Sharu Theresa Jose and Osvaldo Simeone

Title: On the Necessity of Overparameterization in the Quantum Walk Optimization Algorithm
Abstract: Quantum algorithms have emerged as a promising tool to solve combinatorial optimization problems. The quantum walk optimization algorithm is one such variational approach that has recently gained attention. In the broader context of variational quantum algorithms, understanding the expressivity of an ansatz has proven critical for evaluating their performance. A key tool for studying expressivity is through the dimension of the dynamic Lie algebra. In this work, we apply the dynamic Lie algebra framework to the quantum walk optimization algorithm to analyze the role of expressivity in its performance. We derive novel upper bounds on the dynamic Lie algebra dimension for QWOA applied to arbitrary optimization problems. As a direct

implication, we show that solving important problems such as unstructured search and Max-Cut on 2-regular graphs and chains requires a highly overparameterized QWOA circuit.

Authors: Guilherme Adamatti Bridi, Debbie Lim, Lirandë Pira, Raqueline Azevedo Medeiros Santos, Franklin de Lima Marquezino and Soumik Adhikary

Title: Online Learning of Pure States is as Hard as Mixed States

Abstract: Quantum state tomography, the task of learning an unknown quantum state, is a fundamental problem in quantum information. In standard settings, the complexity of this problem depends significantly on the type of quantum state that one is trying to learn, with pure states being substantially easier to learn than general mixed states. A natural question is whether this separation holds for any quantum state learning setting. In this work, we consider the online learning framework and prove the surprising result that learning pure states in this setting is as hard as learning mixed states. More specifically, we show that both classes share almost the same sequential fat-shattering dimension, leading to identical regret scaling. We also generalize previous results on full quantum state tomography in the online setting to the epsilon-realizable setting and to learning the density matrix only partially, using smoothed analysis.

Authors: Maxime Meyer, Soumik Adhikary, Naixu Guo and Patrick Rebentrost

Title: Optimising entanglement distribution policies under classical communication constraints assisted by reinforcement learning

Abstract: Quantum repeaters play a crucial role in the effective distribution of entanglement over long distances. The nearest-future type of quantum repeater requires two operations: entanglement generation across neighbouring repeaters and entanglement swapping to promote short-range entanglement to long-range. For many hardware setups, these actions are probabilistic, leading to longer distribution times and incurred errors. Significant efforts have been vested in finding the optimal entanglement-distribution policy, that is, the protocol specifying when a network node needs to generate or swap entanglement, such that the expected time to distribute long-distance entanglement is minimal. This problem is even more intricate in more realistic scenarios, especially when classical communication delays are taken into account.

In this work, we formulate our problem as a Markov decision problem and use reinforcement learning to optimise over centralised strategies, where one designated node instructs other nodes which actions to perform. Contrary to most RL models, ours can be readily interpreted. Additionally, we introduce and evaluate a fixed local policy, the 'predictive swap-asap' policy, where nodes only coordinate with nearest neighbours. Compared to the straightforward generalization of the common swap-asap policy to the scenario with classical communication effects, the 'wait-for-broadcast swap-asap' policy, both of the aforementioned entanglement-delivery policies are faster at high success probabilities. Our work showcases the merit of considering policies acting with incomplete information in the realistic case when classical communication effects are significant.

Authors: Jan Li, Tim Coopmans, Patrick Emonts, Kenneth Goodenough, Jordi Tura and Evert van Nieuwenburg

Title: Optimization Framework for Data-Adaptive and Hardware-Efficient Quantum Data Embedding

Abstract: Quantum data embedding is essential for applying quantum machine learning algorithms to classical data. This process plays a critical role, as it influences the expressivity and performance

of quantum machine learning models. In this work, we propose a hybrid framework that constructs data-adaptive and hardware-efficient quantum embeddings through a three-stage optimization process. The framework begins by training an autoencoder to extract feature vectors from raw data, followed by a multi-objective genetic algorithm that searches for shallow quantum embedding circuits optimized for classification performance and hardware efficiency. To enhance data adaptability, the circuit parameters are fine-tuned using a neural network-based optimizer. We evaluated our framework on binary classification tasks using the MNIST and CIFAR-10 datasets with a quantum support vector machine. Experimental results show improved classification accuracy over conventional quantum embeddings and several classical kernel SVMs. These findings highlight the potential of hybrid classical-quantum optimization for efficient and effective quantum data embedding on noisy intermediate-scale quantum devices.
Authors: Seokhoon Jeong and Daniel Kyungdeock Park

Title: Optimizing Quantum Time Dynamics with Classical Support

Abstract: Quantum time evolution is a foundational application of quantum computing and is widely regarded as one of the most promising candidates for achieving quantum advantage. The current state-of-the-art approach for simulating the time evolution of a k-local Hamiltonian relies on product formulas, constructed via a Trotterization procedure. Trotter circuits are straightforward to construct and come with rigorous analytical expressions that quantify the error as a function of circuit depth and enable trading one for the other and vice versa. However, achieving high precision often demands a number of operations that are prohibitively large for current quantum hardware. Variations of Trotter methods leading to shallower circuits have, therefore, been studied. Additionally, these methods are mainly suitable for real-time evolution and are non-trivial to apply to imaginary time evolution.

Variational Quantum Time Evolution offers an interesting alternative to Trotter simulation that is directly applicable to both real and imaginary time. It employs McLachlan's variational principle to propagate the parameters of a variational ansatz over discrete time steps. Here, the goal is to follow the exact evolution trajectory as closely as possible. The feasibility of the practical realization of Variational Quantum Time Evolution with a quantum computer strongly depends on the required number of circuit evaluations and measurements. The parameter evolution is defined via a system of linear equations which depends on the Quantum Geometric Tensor and a gradient. Notably, the number of entries in the gradient and Quantum Geometric Tensor scales linearly and quadratically with the number of parameters, respectively. Since the underlying system of linear equations tends to be ill-conditioned, these quantities have to be evaluated with high precision. To suppress shot noise, many measurements have to be taken, and the overall demand on quantum resources is significant.

In this work, we show that the inherent structure in the Quantum Geometric Tensor and gradient entries can allow for high precision evaluation with classical methods, given sufficiently shallow ansatz circuits and short evolution times. Thus, we demonstrate that Variational Quantum Time Evolution can be executed as a true hybrid scheme in which classical resources are employed whenever possible and quantum resources are used where needed, working together towards outperforming purely classical approaches. In certain cases, the propagation of Variational Quantum Time Evolution parameters can be entirely simulated classically, such that available quantum resources can be focused on truly difficult tasks, such as sampling from the final state. This workflow can significantly lower the barrier for practically studying quantum algorithms relying on time-evolution circuits, such as subspace expansions, Gibbs state preparation, and combinatorial optimization on quantum hardware.

Authors: Marc Drudis, Alberto Baiardi, Mattia Chiurco, Francesco Tacchino, Stefan Woerner and Christa Zoufal

Title: Optimizing Shadow Tomography for Many-Body Observables

Abstract: Scalable and robust information extraction is crucial for advancing quantum technologies. Shadow tomography, a powerful method for estimating properties of quantum states, can be implemented using generalized measurements, which offers a more general and conceptually simpler framework than traditional approaches based on random unitaries. This generalization facilitates theoretical analysis, enabling precise characterization of sample complexity and robustness against noise. Additionally, it opens the door to optimizing the measurement basis with respect to a specific set of observables, which is essential for practical applications in quantum simulation and computation. Our optimization scheme leverages convex combinations of positive operator-valued measurements to efficiently generate valid candidate measurements. We demonstrate that multi-qubit measurements are particularly advantageous for the tomography of many-body Hamiltonians, where they outperform single-qubit strategies in terms of efficiency and accuracy. Notably, for the Hydrogen molecule, our method achieves an improved maximum shadow norm compared to previous results, directly enhancing both the robustness and the complexity scaling of the shadow tomography protocol. These advances pave the way for more practical and resilient quantum information extraction in near-term quantum devices.

Authors: Daniel Pranjic and Semih Celiksümer

Title: Parameterized IQP-QCBM generative model: universality with hidden units and kernel-adaptive efficient training on classical hardware

Abstract: In a series of recent works, a promising quantum generative model based on parameterized instantaneous polynomial quantum circuits has emerged. Its unique feature is that it is arguably classically intractable, as these circuits cannot be efficiently sampled from, while, surprisingly, the training can be executed on a classical machine for a particular loss function, namely a squared MMD loss function. However, a few problems limit this model's broader utility. First, the basic model was proven not to be universal for generating arbitrary distributions, although it was suspected that its marginals can be, much like Boltzmann machines achieve universality by utilizing hidden layers. Second, the restriction to the squared MMD loss may be suboptimal in many applications.

Here we provide significant strides in both directions. First, we provide two proofs of the suspected universality, with a main construction that is frugal, requiring only doubling the qubit number. Second, we generalize the possible training methods. The squared MMD loss function depends on a kernel; in previous works, this was the Gaussian kernel. Here we propose an approach to realize task-dependent kernel functions: we use a classical neural network to parameterize the spectral measure of the kernel and use it in the definition of MMD. In addition, we provide a GAN-like training procedure: we tune the hyperparameters of the quantum circuit to minimize the MMD loss while the classical neural network is optimized to challenge the quantum circuit. We believe these two innovations make these models much more usable in practice.

Authors: Andrii Kurkin, Kevin Shen, Susanne Pielawa, Hao Wang and Vedran Dunjko

Title: Partition function estimation with a quantum coin toss

Abstract: Estimating quantum partition functions is a critical task in a variety of fields. However, the problem is classically intractable in general due to the exponential scaling of the Hamiltonian dimension in the number of particles. This paper introduces a quantum algorithm for estimating the partition function of a generic Hamiltonian up to multiplicative error based on a quantum coin toss. The coin is defined by the probability of applying the quantum imaginary-time evolution propagator at inverse temperature to the maximally mixed state, realized by a block-encoding of the propagator into a unitary quantum circuit followed by a post-selection measurement. Our algorithm does not use costly subroutines such as quantum phase estimation or amplitude amplification, and the binary nature of the coin allows us to invoke tools from Bernoulli-process analysis to prove a runtime scaling that is quadratically better than previous general-purpose algorithms using similar quantum resources. Moreover, since the coin is defined by a single observable, the method lends itself well to quantum error mitigation. We test this in practice with a proof-of-concept nine-qubit experiment, where we successfully mitigate errors through a simple noise-extrapolation procedure..

Authors: Thais Lima Silva, Lucas Borges and Leandro Aolita

Title: Physics-Informed Neural Networks for Simulating Open Quantum Systems

Abstract: In the rapidly evolving field of quantum computing, the study of open quantum systems remains crucial. Understanding the dynamics of these systems drives the development of better quantum devices, error mitigation, and error correction strategies. These systems are commonly described by a density matrix evolving under the Gorini-Kossakowski-Sudarshan-Lindblad master equation. An open quantum system can be represented by a density matrix that evolves according to the Gorini-Kossakowski-Sudarshan-Lindblad master equation. When working with continuous-variable open quantum systems, it is beneficial to transform the master equation into a partial differential equation describing the evolution of a quasi-probability distribution instead. In this work, we make use of the Husimi Q-Function representation that evolves according to Fokker-Planck equations. Using this representation, we are able to simulate open quantum systems using state-of-the-art physics-informed neural networks (PINNs). PINNs provide a powerful alternative to traditional numerical approaches for solving differential equations such as finite difference and finite element schemes. Unlike these traditional approaches, PINNs do not require costly evaluations on fine grids and can incorporate experimental data. In this work, we train PINNs to solve several equations that govern the dynamics of the Husimi Q-Function. We evaluate the performance of several architectures and loss functions for PINNs and compare the obtained solutions to analytical solutions where possible. We further compare several numerical integration techniques for estimating key observables, such as the average photon number and average displacement, from the PINN solutions. Each technique is evaluated based on how well the estimated observables align with reference values, obtained via numerical integration of analytical solutions, or how well the estimated observables reproduce expected trends. Overall, our results establish physics-informed neural networks as a promising machine learning-based framework for simulating open quantum systems.

References

1. V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, "Completely positive dynamical semigroups of n-level systems," Journal of Mathematical Physics, vol. 17, no. 5, pp. 821–825, 1976.

2. G. Lindblad, “On the generators of quantum dynamical semigroups,” Communications in Mathematical Physics, vol. 48, pp. 119–130, 1976.

3. H. J. Carmichael, Statistical methods in quantum optics 1: master equations and Fokker-Planck equations. Springer Science & Business Media, 2013.

4. I. E. Lagaris, A. Likas, and D. I. Fotiadis, “Artificial neural networks for solving ordinary and partial differential equations,” IEEE transactions on neural networks, vol. 9, no. 5, pp. 987–1000, 1998.

5. M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” Journal of Computational Physics, vol. 378, pp. 686–707, 2019.

Authors: Shivani Pillay, Ilya Sinayskiy and Francesco Petruccione

Title: Pitfalls in the hunt for scalable parameterized quantum models

Abstract: Identifying circuit architectures and optimization strategies that are free from exponential concentration is a core quest in the hunt for scalable models in variational quantum computing. So far, there is an increasingly large number of proposals for circumventing exponential concentration—both those explicitly claiming to avoid or mitigate barren plateaus, including special circuit architectures and alternative initialization strategies, as well as those mentioned informally, such as sample-based optimization or quantum natural gradient descent. Here we argue that, given the subtle interplay between quantum measurements and classical processing strategies, care needs to be taken to determine whether these approaches do in fact help in practice. In particular, any procedures in variational quantum computing involve estimating some variable-dependent quantities. We show that from the practical perspective, if outcome probabilities exponentially concentrate, the measurement outcomes together with post-processing contain no information about the variables. Based on these results, we provide a practical step-by-step guideline for identifying whether a given procedure can circumvent exponential concentration. This guideline can be used to debunk some previously considered barren-plateau-free methods, including natural gradient descents, sample-based CVaR optimization, agnostic classical neural network-assisted initialization, and rescaled gradient approaches.

Authors: Reyhaneh Aghaei Saem, Behrang Tafreshi, Zoe Holmes and Supanut Thanasilp

Title: Predictive Control with Hybrid Depth-Infused Quantum Neural Networks

Abstract: Accurate prediction and stabilization of blast furnace temperatures are crucial for optimizing the efficiency and productivity of steel production. Traditional methods often struggle with the complex and non-linear nature of the temperature fluctuations within blast furnaces. This paper proposes a novel approach that combines hybrid quantum machine learning with pulverized coal injection control to address these challenges. By integrating classical machine learning techniques with quantum computing algorithms, we aim to enhance predictive accuracy and achieve more stable temperature control. For this, we utilized a unique prediction-based optimization method. Our method leverages quantum-enhanced feature space exploration and the robustness of classical regression models to forecast temperature variations and optimize pulverized coal injection values. Our results demonstrate a significant improvement in prediction

accuracy of over 25 percent, and our solution improved temperature stability to 7.6 degrees of the target range from the earlier variance of 50 degrees, highlighting the potential of hybrid quantum machine learning models in industrial steel production applications.
Authors: Nayoung Lee, Minsoo Shin, Asel Sagingalieva, Ayush Joshi Tripathi, Karan Pinto and Alexey Melnikov

Title: Probabilistic Greedy Behaviour of the Equivariant Quantum Circuit

Abstract: In this paper, we present experimental evidence that the Equivariant Quantum Circuit (EQC) to solve the TSP performs no better than a classical Probabilistic Nearest Neighbour (PNN) Algorithm. The original EQC work reported near-optimal performance on TSP instances with 20 nodes and fewer. Despite strong interest since its 2023 publication, we find that the EQC's performance at Depth 1 is statistically indistinguishable from a PNN baseline in terms of optimality gaps on both TSP instances of uniform node locations, as well as TSPLIB instances between 5 and 55 nodes. On a set of handcrafted adversarial TSP instances designed to expose local decision making, the tours produced by the EQC at depths 1 to 4 are largely similar to the tours produced by the PNN baseline. Lastly, we evaluated a classical model Structure2Vec (S2V), and found this model consistently performs better than the PNN baseline and the EQC.
Authors: Jonathan Teo, Xin Wei Lee and Hoong Chuin Lau

Title: Problem-informed Graphical Quantum Generative Learning

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

Title: Prospects for quantum advantage in ML from the representability of quantum functions

Abstract: Quantum Machine Learning (QML) offers a promising avenue for computational advantage, but identifying which problems and models can outperform classical methods remains a central challenge. A critical gap exists in systematically connecting the structural properties of Parameterized Quantum Circuits (PQCs) (such as gate sets, depth, and architecture) to the mathematical nature of the functions they can represent and their susceptibility to classical simulation. This work introduces a unifying framework to bridge this gap. We classify PQCs into three nested classes based on the classical representability of the functions they generate: (1) 'Identifiable & Evaluatable' circuits, which are fully and efficiently simulable; (2) 'Evaluatable'

circuits, whose output functions belong to an efficient classical family but may be hard to identify for a given set of parameters; and (3) 'Quantum Functions', which may not admit an efficient classical description. By mapping concrete PQC architectural and resource constraints to this classification, we demonstrate that many known simulable circuits (e.g., those with logarithmic depth, low non-Clifford gate counts, or free-fermionic dynamics) fall into Class 1. We find that 'flipped architectures' with specific properties like low-doping encodings are primary examples of Class 2, making them vulnerable to classical surrogate models. Our analysis reveals that the most promising candidates for quantum advantage are Class 3 circuits, which typically involve deep and highly entangling structures that evade known dequantization strategies. Ultimately, this framework provides a structured lens to assess QML models, guiding the design of circuits with a greater potential for demonstrating quantum advantage.
Authors: Carlos Bravo-Prieto, Elies Gil-Fuster, Sergi Masot-Llima and Tommaso Guaita

Title: Provably Robust Training of Quantum Circuit Classifiers Against Parameter Noise

Abstract: Advancements in quantum computing have spurred significant interest in harnessing its potential for speedups over classical systems. However, noise remains a major obstacle to achieving reliable quantum algorithms. In this work, we present a provably noise-resilient training theory and algorithm to enhance the robustness of parameterized quantum circuit classifiers. Our method, with a natural connection to Evolutionary Strategies, guarantees resilience to parameter noise with minimal adjustments to commonly used optimization algorithms. Our approach is function-agnostic and adaptable to various quantum circuits, successfully demonstrated in quantum phase classification tasks. By developing provably guaranteed optimization theory with quantum circuits, our work opens new avenues for practical, robust applications of near-term quantum computers.
Authors: Lucas Tecot, Di Luo and Cho-Jui Hsieh

Title: QAOA based Neural Architecture Search

Abstract: Neural Architecture Search (NAS) automates neural network (NN) design but struggles with the exponential growth of the search space (SSp) as layers and choices increase. Quantum computing offers a potential solution. This paper investigates using the Quantum Approximate Optimization Algorithm (QAOA) to enhance NAS by encoding the Mean Squared Error (MSE) loss as a quantum Hamiltonian. We express architecture selection via binary variables mapped to Pauli-Z operators but identify a key obstacle: nonlinear operations (e.g., ReLU, pooling) prevent separation of classical coefficients and quantum observables, rendering the Hamiltonian incompatible with QAOA. Our results highlight fundamental limits and motivate alternative quantum-classical strategies.
Authors: Alvaro Romero Mato, Boyang Chen, Nathan Eskue and Vahid Nasrabadi

Title: QCirCNN: a photonic-native quantum circular convolution network based on circulant matrices

Abstract: Convolutional neural networks have revolutionised the field of artificial intelligence, especially in computer vision, where they outperformed multilayer perceptrons. Although quantum convolutional neural networks (QCNNs) have been widely studied, implementing convolutional operations on quantum computers represents a challenging task. Here, we propose a quantum circular convolutional neural networks (QCirCNN) based on circulant matrices, which can be implemented natively on photonic devices using linear optical interferometers and adaptive

measurements. We trained this model on the MNIST classification dataset, and observed accuracies outperforming classical CNNs, while having almost half the number of trainable parameters. Moreover, our architecture is resilient to barren plateaus due to its Hamming-weight preserving design. These results suggest that QCirCNN offers a scalable and hardware-aligned approach to quantum vision tasks.
Authors: Daphne Wang, Anthony Walsh, Hugo Fruchet, Ariane Soret and Pierre-Emmanuel Emeriau

Title: Quantum Algorithm for Solving Nonlinear Differential Equations Based on Physics-Informed Effective Hamiltonians
Abstract: We propose a distinct approach to solving linear and nonlinear differential equations on quantum computers by encoding the problem into the ground states of effective Hamiltonian operators. Our algorithm relies on constructing such operators in the Chebyshev space, where an effective Hamiltonian is a sum of global differential and data constraints. Once the effective Hamiltonian is formed, solutions of differential equations can be obtained using the ground state preparation techniques such as quantum imaginary-time evolution based on quantum singular value transformation, thus bypassing variational search. Unlike approaches based on discrete grids, the algorithm enables evaluation of solutions beyond fixed grid points and implements constraints in a physics-informed and data-driven way. Our proposal inherits the best traits from quantum machine learning-based differential equation solvers and quantum linear algebra-based approaches, offering a robust strategy for quantum scientific computing in the early fault-tolerant era.
Authors: Hsin-Yu Wu, Annie E. Paine, Evan Philip, Antonio A. Gentile and Oleksandr Kyriienko

Title: Quantum algorithms for representation-theoretic multiplicities
Abstract: Kostka, Littlewood-Richardson, Plethysm and Kronecker coefficients are the multiplicities of irreducible representations in the decomposition of representations of the symmetric group that play an important role in representation theory, geometric complexity and algebraic combinatorics. We give quantum algorithms for computing these coefficients whenever the ratio of dimensions of the representations is polynomial and study the computational complexity of this problem. We show that there is an efficient classical algorithm for computing the Kostka numbers under this restriction and conjecture the existence of an analogous algorithm for the Littlewood-Richardson coefficients. We argue why such classical algorithm does not straightforwardly work for the Plethysm and Kronecker coefficients and conjecture that our quantum algorithms lead to superpolynomial speedups for these problems. The conjecture about Kronecker coefficients was disproved by Greta Panova in arXiv:2502.20253 with a classical solution which, if optimal, points to a O(n^{4+2k}) vs Omega(n^{4k^2+1}) polynomial gap in quantum vs classical computational complexity for a free integer parameter k.
Authors: Martin Larocca and Vojtech Havlicek

Title: Quantum Chebyshev Probabilistic Models for Fragmentation Functions
Abstract: We propose a quantum protocol for efficiently learning and sampling multivariate probability distributions that commonly appear in high-energy physics. Our approach introduces a bivariate probabilistic model based on generalized Chebyshev polynomials, which is pretrained as an explicit circuit-based model for two correlated variables, and sampled efficiently with the use of quantum Chebyshev transforms. As a key application, we study the fragmentation functions of

charged pions and kaons from single-inclusive hadron production in electron-positron annihilation. We learn the joint distribution for the momentum fraction z and energy scale Q in several fragmentation processes. Using the trained model, we infer the correlations between z and Q from the entanglement of the probabilistic model, noting that the developed energy-momentum correlations improve model performance. Furthermore, utilizing the generalization capabilities of the quantum Chebyshev model and extended register architecture, we perform a fine-grid multivariate sampling relevant for FF dataset augmentation. Our results highlight the growing potential of quantum generative modeling for addressing problems in scientific discovery and advancing data analysis in high-energy physics.
Authors: Jorge Martínez de Lejarza, Hsin-Yu Wu, Oleksandr Kyriienko, German Rodrigo and Michele Grossi

Title: Quantum convolutional neural networks produce higher variance in regression tasks
Abstract: Being a specific type of parametrized quantum circuits, quantum convolutional neural networks (QCNNs) is a widely used tool for performing quantum machine learning tasks on labeled quantum states. Such circuits have layered structure, and after each layer a subset of qubits of the processed state is measured or traced out. At the end of the network, one typically measures a local observable. In our work, we demonstrate that if such tools are applied to solving regression tasks on labeled quantum data, it generally results in larger label prediction variance. We show that reason for this is, essentially, the number of distinct eigenvalues of the observable one measures after the application of a QCNN or other variational ansatz.
Authors: Andrey Kardashin, Vladimir Palyulin and Konstantin Antipin

Title: Quantum Differential Privacy in Quantum Federated Learning
Abstract: Quantum federated learning (QFL), which has potential to enable distributed quantum machine learning (QML), is exposed to various privacy threats exploiting the QFL's structural property. To deal with the privacy leakage of the QFL, this paper proposes novel QFL structure adopting quantum teleportation (QT) and quantum differential privacy (QDP), which considers communication efficiency, privacy, and training accuracy at the same time. The effectiveness of the proposed QDP has been confirmed through comprehensive experiments, which show a notable reduction in attack success rates and improvement of communication efficiency.
Authors: Chaemoon Im and Joongheon Kim

Title: Quantum Extreme Learning Machines: Insights from Industrial and Real-World Applications
Abstract: This talk will focus on the application of Quantum Artificial Intelligence (QAI), specifically a quantum machine learning algorithm called Quantum Extreme Learning Machines (QELMs), to address three industrial and real-world societal problems. QELMs implement quantum mechanical principles to enable the efficient training of simple classical machine learning models (e.g., linear regression), even with a few features. We studied one application of industrial elevators built by Orona, Spain, whereas the other two applications come from real-world societal contexts in Norway: the cancer registry system from the Cancer Registry of Norway and IoT-based healthcare services from Oslo City. In all three cases, the context was regression testing of classical software implemented in these applications with QELMs. We assessed QELMs both with and without noise. Our results showed that their performance without noise was at least similar to, or better than, classical machine learning, whereas their performance was affected by noise. Including noise during machine learning training and testing improved their performance. Moreover, using error

mitigation methods such as IBM's ZNE method further enhanced performance. In general, our results demonstrate their potential in real-world applications; however, maximizing their practical benefit requires developing new methods to handle noise.
Authors: Shaukat Ali

Title: Quantum generative diffusion models for medical imaging

Abstract: In the field of Quantum Generative Artificial Intelligence (QGenAI), Quantum Diffusion Models (QDMs) are an emerging class of algorithms that aim to leverage quantum mechanical dynamics to enhance the performance of their classical counterparts. The overall workflow of QDMs follows the same principles as classical diffusion models (DMs), comprising a forward (diffusion) process and a backward (denoising) process. In the forward process, an N-qubit quantum state, represented by a density matrix, is evolved over time into a progressively noisier state, until it reaches a maximally mixed state. Conversely, the backward process employs a parametrized learning model to reconstruct the less noisy state at $t-1$ from the one at time t, in a similar way as in classical DMs.
In our recent work, we introduced two scalable, physics-inspired QDM protocols for image generation. In the first approach, we demonstrated that a hybrid quantum-classical stochastic forward dynamic yields statistically more robust generative models, achieving lower Frechet Inception Distance (FID) scores when generating MNIST images compared to models based solely on classical or total quantum dynamics. Using the quantum stochastic walks (QSWs) approach to model the diffusion, each image pixel is treated as an independent quantum walker moving on a graph with nodes corresponding to the intensity levels. The dynamics are governed by the Kossakowski–Lindblad–Gorini master equation. The backward process then employs an artificial neural network to learn and recover the original data distribution. Additionally, we proposed an algorithm implementable on real quantum hardware that utilizes the intrinsic noise of quantum devices to drive the generation of synthetic data.
The principal contribution of this project is to extend our previous QDMs to the domain of real-world data, particularly in medical imaging, where synthetic data generation is critical due to the limited availability of annotated records and the need to preserve patient privacy. Specifically, we investigate the behavior and convergence of the diffusion dynamics under varying image sizes and discrete intensity levels. Furthermore, we analyze the incorporation of multiple independent learning models and composite loss functions, such as cross-entropy for pixel-level intensity prediction and structural similarity metrics for full-image comparison.
In this work, we apply the QSW framework to generate BloodMNIST images from the MedMNIST dataset. Each sample comprises 64×64 pixels with 32 discrete grayscale intensity levels. Our initial experiments indicate that it is indeed possible to recover image data from a quantum stochastic diffusion process, although this requires a significant effort in the design and training of the backward model.
Authors: Francesco Aldo Venturelli, Marco Parigi, Stefano Martina, Natalia Muñoz Moruno, Filippo Caruso, Alba Cervera Lierta and Miguel Ángel González Ballester

Title: Quantum Graph Neural Networks for the Travelling Salesman Problem

Abstract: Solving combinatorial optimization problems effectively and more efficiently with quantum computers can have huge impact in many areas of industry. Combining traditional solvers with learned based methods, such as neural combinatorial optimization (NCO) can achieve the best of both methods. In this work, we propose a Quantum Graph Neural Network architecture which is tailored to problems such as TSP, but yet is general enough to be applicable to other graph

learning problems. The model uses subspace preserving circuits, which have demonstrated promising features in quantum learning problems, to encode and transform graph data in fixed Hamming-weight subspaces. Using these circuits to generate candidate heatmaps containing candidate TSP tours, we can improve ultimate solutions found by traditional classical optimization solvers.
Authors: Snehal Raj, Brian Coyle, Léo Monbroussou and Elham Kashefi

Title: Quantum Hyperdimensional Computing for Pattern Completion
Abstract: The pattern completion problem is a core task in computational learning and symbolic reasoning, with applications in error correction, language processing, and beyond. In its basic form, the problem asks: given an incomplete or noisy input, can we reconstruct the complete pattern? In the binary case, this involves a partially observed binary sequence y of length n minus k, and an unknown suffix omega of length k, with the goal of producing a completed sequence z = y concatenated with omega, subject to a set of constraints C. These constraints may represent consistency checks, logical rules, optimization criteria, or problem-specific formulas. In the simplest case, if C requires the completed sequence to belong to a regular language L, the task reduces to finding omega such that y omega belongs to L. When k is zero, this becomes a membership problem, which is undecidable in general for Turing machines. In this work, we explore connections with automata theory, hyperdimensional computing (HDC), and quantum computing, aiming to solve completion tasks in polynomial time using recent quantum HDC architectures. Unlike classical approaches, these architectures avoid strong assumptions from the Chomsky hierarchy and offer approximate, heuristic solvers with better guarantees.
Authors: Leonardo Lavagna, Francesca De Falco and Massimo Panella

Title: Quantum medical image encoding and compression using Fourier-based methods
Abstract: Quantum image processing (QIMP) is a growing field within quantum computing applications, aiming to offer computational advantages over classical image processing methods. In most QIMP algorithms, the first critical step is to encode classical image information into a quantum circuit. However, most existing quantum image encoding methods—based on either the Flexible Representation of Quantum Images (FRQI) or the Novel Enhanced Quantum Representation (NEQR)—require a number of quantum gates nearly twice the number of pixels in the image. As a result, simulating even a modest-sized image (e.g., 1024 × 1024) becomes computationally demanding. In this work, we propose a quantum image encoding method that significantly reduces the number of gates compared to existing approaches. Unlike conventional methods, our approach exploits the effectiveness of the discrete Fourier transform (DFT) for image data compression. After compressing the image, we employ the Fourier-series loader circuit to encode the compressed image into a quantum circuit. This procedure yields an efficient quantum circuit, particularly when the image is highly compressible via DFT. We demonstrate our method using various high-resolution (1024 × 1024) medical images captured during Bilateral Axillo-Breast Approach (BABA) robotic thyroidectomy surgeries. Our results show that the proposed method achieves approximately a 98% reduction in gate count compared to existing methods such as FRQI and NEQR—a significant improvement. Furthermore, we introduce two additional compression techniques to further reduce the number of gates and preprocessing time, with negligible loss in image quality. We propose our image encoding strategy as a valuable option for large-scale medical imaging applications.
Authors: Taehee Ko, Hyeong Won Yu, Inho Lee, Sangkook Choi and Hyowon Park

Title: Quantum Neural Density Functionals in Density Functional Theory

Abstract: Quantum computing holds promise for accurately simulating electronic structures but faces practical constraints due to high resource requirements. Density functional theory, while foundational in quantum chemistry, often struggles to capture strongly correlated systems due to inherent approximations in exchange-correlation functionals. Here, we introduce Quantum Neural Functionals, a hybrid quantum-classical framework leveraging quantum neural networks to enhance density functional theory. Motivated by the famous Hohenberg-Kohn and Levy-Lieb theorems, our approach explicitly encodes quantum many-body correlations directly from classical charge densities. Crucially, our quantum neural network architecture incorporates rigorous symmetry constraints ensuring rotational invariance, significantly improving model trainability, mitigating optimization issues such as barren plateaus, and enhancing generalization. Rigorous testing on challenging cases like H2 dissociation reveals substantial performance gains over classical and naive quantum methods, highlighting the potential of symmetry-informed Quantum Neural Functionals in advancing density functional theory.

Authors: Minh Triet Chau, Hyeokjea Kwon, Sung Won Yun, Kevin Ferreira, Thi Ha Kyaw and Jack Baker

Title: Quantum Neural Networks Facilitating Quantum State Classification

Abstract: Entangled quantum states exhibit non-local correlations that defy classical notions of locality and serve as essential resources in quantum protocols. However, classifying quantum states—especially in multi-qubit systems where the number of potential subclasses increases with the number of qubits—remains a significant challenge. In this study, we propose an approach that classifies n-qubit quantum states using only n qubits while distinguishing between various subclasses, elevating the use of quantum neural networks. For this, instead of employing a traditional feature mapping circuit, we integrate the ansatz of quantum neural networks with a problem-inspired circuit. The problem-inspired circuit is equipped with parametrized two-qubit unitary operators, constructed using Sz.-Nagy's dilation theorem. This resource-efficient approach generates various classes of quantum states by varying the entangling power of the resulting global unitary operator. We also visualize and quantify the mitigation of barren plateaus, demonstrating improved trainability and expressivity of the proposed ansatz. The designed quantum neural network demonstrates efficiency in binary and multi-class classification tasks. This work establishes a foundation for classifying multi-qubit quantum states with remarkable accuracy. Notably, the proposed architecture significantly reduces quantum resource demands by only utilizing the number of qubits equal to that present in the input state, with strong classification performance.

Authors: Diksha Sharma, Vivek Balasaheb Sabale, Atul Kumar and Thirumalai M.

Title: Quantum Optimization Towards Large-Scale Molecular Docking on a Quantum Computer

Abstract: Molecular docking is a foundational computational task in drug discovery, wherein the objective is to efficiently identify optimal binding poses between a ligand and a target receptor protein. Due to the combinatorial explosion of possible binding configurations, molecular docking remains a computationally intensive problem, especially at scale. Recent work reformulates docking as a Max-Weighted Clique problem on a compatibility graph, where nodes correspond to candidate fragment alignments and edge weights encode both spatial and physicochemical compatibility, derived from precomputed molecular descriptors and experimental binding data. In

this work, we present a quantum-classical hybrid approach for molecular docking leveraging the Max-Weighted Clique formalism with a novel Variational Multibasis Encoding strategy, which enables efficient encoding of classical binary variables with Bloch sphere vectors. The molecular docking problem is mapped to a cost Hamiltonian that is minimized within a variational framework, optimized via gradient-based techniques. Our results highlight the feasibility and scalability of quantum-enhanced optimization for large-scale structure-based drug design and point towards the broader utility of advanced encoding techniques in quantum optimization for computational biology.
Authors: Tianqi Chen and Jian Feng Kong

Title: Quantum Principal Basis Learning (qPBL) for image classification
Abstract: This work presents preliminary results of a quantum machine learning framework designed to learn an optimal basis transformation such that the principal components of the superposition of images with similar features converge to a common eigenstate. The framework aims to identify a transformation that emphasizes the key features shared across all superpositions with the same label. The algorithm searches over a subspace of SU(2l), learning a sequence of unitary operations that aligns the superpositions into a shared eigenspace in the principal component basis.
Authors: Gabriel Mejia Ruiz, Eileen Kuhn and Achim Streit

Title: Quantum reservoir computing with a single quantum chaotic node
Abstract: Quantum reservoir computing (QRC) is an emerging paradigm that employs quantum dynamical systems as reservoir for machine learning tasks. Due to quantum superposition and entanglement, QRC is expected to outperform classical reservoir computing. Most earlier works on QRC have utilised extended quantum system as reservoir. In this work, we create a framework for QRC using a single quantum chaotic map and perform many benchmark tests for prediction and memory capacity. In particular, we demonstrate entanglement classification using quantum data, and also chaotic time series prediction using classical data. Our results show that quantum chaos aids certain learning tasks with QRC.
Authors: Santhanam Madabushi Srinivasan and Nisarg Vyas

Title: Quantum Scalar Field Theoretic Extension of Boltzmann Machines to Solve a Class of Moment Matching Problems
Abstract: The Boltzmann machine is a machine learning model originated from the toy model of magnetic materials in statistical mechanics. It can approximate a probability distribution by adjusting the set of potential parameters and the number of units. In particular, over the last decade, there has been a significant amount of research on approximating ground state wave functions of quantum many-body systems via the Boltzmann machine. However, for Boltzmann machines to achieve high expressive power, increasing the number of units is unavoidable. The computational complexity in this case is known to be NP-hard, hence we have to develop alternative methods to achieve scalability. Notably, it is a significant challenge to improve the representativity of the model enough to analyzing quantum phenomena, without increasing the number of units. In this study, firstly, we introduce the Scalar Field Machine (SFM) as a generalized model of the Boltzmann machine, originated from phi4 scalar field model in constructive quantum field theory. By utilizing it, we show some moment matching problems can be solved without increasing the number of units. As an application, we demonstrate that short-time entangled

behavior of the dynamically decoupling quantum harmonic oscillators can be approximated by the SFM. The dynamics is constructed via the stochastic quantization, which is equivalent to the canonical quantization. However, for long-time dynamics, the SFM approximation begins to break down, hence it needs to update the distribution successively. The optimal update rule for this is currently under investigation.
Authors: Takahiro Kajisa

Title: Quantum simulation in the Heisenberg picture via Vectorization

Abstract: A central challenge in quantum many-body physics is to understand the behavior of operators under time evolution in the Heisenberg picture. In this work, we propose a framework to perform quantum simulation in the Heisenberg picture using quantum computers, achieved by encoding Heisenberg operators as quantum states in a doubled Hilbert space through the vectorization map. This enables the execution of a multitude of useful subroutines and algorithms, including directly sampling from the Pauli distribution of an operator, estimation of physically interesting quantities such as two-point correlators, operator stabilizer entropies, entanglement entropies, and the statistical moments of superoperators over Heisenberg operators. To demonstrate the practical utility of our framework, we describe a proposal to probe the spreading of quantum information due to a 2D lattice Hamiltonian, and estimate the resources required to implement this proposal on a quantum computer featuring a square grid topology, such as existing ones based on superconducting qubits.
Authors: Shao Hen Chiew, Armando Angrisani, Zoë Holmes and Giuseppe Carleo

Title: Quantum Spectral Clustering: Comparing Parameterized and Neuromorphic Quantum Kernels

Abstract: We undertake a comprehensive comparison between two quantum-inspired kernel methods: a parameterized quantum fidelity kernel and a quantum leaky integrate-and-fire neuromorphic kernel, within the context of spectral clustering. As a baseline, we also include the classical radial basis function kernel. Our goal is to assess how these distinct quantum data-encoding and similarity-measurement strategies perform on both low- and high-dimensional datasets, and to understand their respective computational trade-offs.

In the parameterized quantum fidelity kernel approach, each d-dimensional feature vector is mapped onto a single qubit via angle encoding on the Bloch sphere, with rotation angles scaled by a set of tunable parameters. The kernel value between two vectors is given by the squared overlap (fidelity) of their parametrically rotated states, as measured by a projector onto the reference state. We optimize the rotation-angle parameters through a grid search that maximizes kernel-target alignment, ensuring that the resulting Gram matrix faithfully reflects pairwise distances in the original feature space.

By contrast, the quantum leaky integrate-and-fire neuromorphic kernel employs population coding to convert real-valued features into ensembles of spike trains. We then compute pairwise similarities using established temporal distance metrics, either the Victor–Purpura or van Rossum metric, which embed the data in a kernel matrix that captures both timing and memory effects. This kernel is subsequently fed into a standard spectral clustering pipeline, enabling a direct performance comparison against the fidelity and radial basis function methods.

On benchmark datasets including synthetic Blobs, Moons, Circles, the Iris dataset, and a pre-processed Sloan Digital Sky Survey catalog, we evaluate clustering quality via label-based metrics and determine the optimal number of clusters using an elbow-style criterion. The neuromorphic kernel consistently outperforms the fidelity and radial basis function kernels on

low-dimensional, non-linearly separable tasks, owing to its intrinsic temporal encoding and adaptive memory. Conversely, the fidelity kernel exhibits superior clustering and more favorable scaling on the higher-dimensional Sloan Digital Sky Survey data. Runtime analysis reveals that the neuromorphic kernel's computational cost grows more steeply with dimension, whereas the fidelity kernel benefits from compact state representations.
These findings highlight complementary regimes of applicability: neuromorphic kernels excel in capturing complex temporal and non-linear structures in small to moderate dimensions, while fidelity-based kernels provide improved performance in large-scale, high-dimensional settings. This work thus charts a principled path for quantum machine learning, suggesting that hybrid or task-specific selection of quantum and neuromorphic kernels can yield significant gains in clustering and beyond.
Authors: Donovan Slabbert, Dean Brand and Francesco Petruccione

Title: Quantum spectral operator learning for solving partial differential equations

Abstract: Solving partial differential equations is computationally expensive, and machine learning methods often depend on large-scale supervised data. While unsupervised frameworks like the Unsupervised Legendre-Galerkin Network mitigate this data dependency, they still face scalability challenges. To overcome these limitations, we propose a quantum-classical hybrid framework for unsupervised spectral operator learning. Our approach first discretizes the partial differential equation into a linear system using the Legendre-Galerkin method. A classical neural network then learns to map the equation's forcing function directly to the variational parameters of a quantum circuit based on the Variational Quantum Linear Solver. This circuit prepares the quantum state encoding the solution's coefficients. We validated our framework on the one-dimensional Helmholtz equation, where the model generalized effectively across 200 unseen test samples, achieving an average relative L2 error below 0.006. By eliminating the data-generation bottleneck and enabling generalization across partial differential equation instances without re-optimization, our work presents a practical pathway for using quantum computers to solve partial differential equations.
Authors: Myeonghwan Seong, Yujin Kim, Chanyoung Kim, Daniel K. Park and Youngjoon Hong

Title: Quantum vs. classical: A comprehensive benchmark study for predicting time series with variational quantum machine learning

Abstract: Variational quantum machine learning algorithms have attracted attention as potential candidates for time series forecasting, with the promise of capturing complex temporal patterns beyond the reach of classical models. Yet, their practical advantage over established classical methods remains uncertain. In this work, we conduct a rigorous and large-scale benchmark comparing several variational quantum algorithms and classical machine learning models on a diverse set of 27 forecasting tasks derived from three chaotic systems. To ensure a fair and meaningful evaluation, all models undergo extensive hyperparameter optimization under comparable constraints. Our results reveal that, despite their theoretical appeal, quantum models often fall short of matching even relatively simple classical baselines in predictive accuracy. By further analyzing how performance scales with model complexity, we gain deeper insight into the capabilities of classical and quantum models. Overall, this study establishes a foundation for guiding future developments in quantum forecasting models.
Authors: Tobias Fellner, David Kreplin, Samuel Tovey and Holm Christian

Title: Quantum-Inspired Self-Attention in a Large Language Model

Abstract: Recent advances in Natural Language Processing (NLP) have been predominantly driven by transformer-based architectures, which rely heavily on self-attention mechanisms to model relationships between tokens in a sequence. Similarly, the field of Quantum Natural Language Processing (QNLP), which seeks to leverage quantum principles to address challenges in language understanding and generation tasks, has seen the recent development of quantum self-attention mechanisms.

We propose a novel quantum-inspired self-attention (QISA) mechanism and integrate it into the full autoregressive language modeling pipeline of GPT-1. To the best of our knowledge, this is the first integration of such kind, as previous quantum self-attention mechanisms have been tested exclusively on text classification. In our experiments, QISA achieves approximately 10.6 times lower cross-entropy loss compared to standard self-attention, while requiring fewer parameters, and only a 2.1 times longer inference time. We provide open-source repository written on the PyTorch + TorchQuantum frameworks:
https://github.com/Nikait/QISA.

Authors: Nikita Kuznetsov and Ernesto Campos

Title: Reducing Circuit Depth of Amplitude Encoding for Gravitational Waves

Abstract: Quantum State Preparation is an increasingly important part of quantum computing and ensuring it can be performed efficiently is essential for the future of quantum algorithms. Here, we look at improving the efficiency of amplitude encoding for numerical and analytical functions. By tolerating a small error in the Grover-Rudolph algorithm, we can greatly reduce the number of gates needed as the number of qubits grows. Previous work has proposed a reduction in the total number of gates needed by replacing controlled rotational gates with fixed rotations for higher order qubits. We improve on this by asymmetrically discretising our distribution, prioritising areas with more variation so that we can more uniformly map the number of gates used to the amount of information in that section. We demonstrate our method by encoding distributions corresponding to gravitational wave template waveforms. We benchmark against other state preparation methods, comparing the total number of controlled rotations needed, and fidelity achieved in each case.

Authors: Elizabeth Sarell, Ashwin Girish, Hector Spencer-Wood, Michael Puerrer, Christopher Messenger, Fiona Speirits and Sarah Croke

Title: Resting-state fMRI Analysis using Quantum Time-series Transformer

Abstract: Resting-state functional magnetic resonance imaging (fMRI) has emerged as a pivotal tool for revealing intrinsic brain network connectivity and identifying neural biomarkers of neuropsychiatric conditions. However, classical self-attention transformer models—despite their formidable representational power—struggle with quadratic complexity, large parameter counts, and substantial data requirements. To address these barriers, we introduce a Quantum Time-series Transformer, a novel quantum-enhanced transformer architecture leveraging Linear Combination of Unitaries and Quantum Singular Value Transformation. Unlike classical transformers, Quantum Time-series Transformer operates with polylogarithmic computational complexity, markedly reducing training overhead and enabling robust performance even with fewer parameters and limited sample sizes. Empirical evaluation on the largest-scale fMRI datasets from the Adolescent Brain Cognitive Development Study and the UK Biobank demonstrates that Quantum Time-series Transformer achieves comparable or superior predictive performance compared to state-of-the-art classical transformer models, with especially pronounced gains in small-sample scenarios.

Interpretability analyses using SHapley Additive exPlanations further reveal that Quantum Time-series Transformer reliably identifies clinically meaningful neural biomarkers of attention-deficit/hyperactivity disorder (ADHD). These findings underscore the promise of quantum-enhanced transformers in advancing computational neuroscience by more efficiently modeling complex spatio-temporal dynamics and improving clinical interpretability.
Authors: Junghoon Justin Park, Jungwoo Seo, Sangyoon Bae, Samuel Yen-Chi Chen, Huan-Hsin Tseng, Jiook Cha and Shinjae Yoo

Title: Retrodictive Approach to Quantum State Smoothing
Abstract: Smoothing is a technique for estimating the state of an imperfectly monitored open system by combining both prior and posterior measurement information. In the quantum regime, current approaches to smoothing either give unphysical outcomes, due to the non-commutativity of the measurements at different times, or require assumptions about how the environment is measuring the system, which with current technology is unverifiable. We propose a novel definition of the smoothed quantum state based on quantum Bayesian retrodiction, which mirrors the classical retrodictive approach to smoothing. This approach always yields physical results and does not require any assumption on the environment. We show that this smoothed state has, on average, greater purity than the state reconstructed using just the prior information. Finally, we make a connection with the well-studied smoothing theory developed by Wiseman.
Authors: Mingxuan Liu, Valerio Scarani, Alexia Auffeves and Kiarn Laverick

Title: Robust and efficient verification of measurement-based quantum computation
Abstract: To achieve reliable measurement-based quantum computation, it is crucial to verify whether the resource graph states are accurately prepared in the adversarial scenario. Previous verification protocols for this task are resource consuming or noise susceptible. Here, we propose a robust and efficient protocol for verifying arbitrary graph states with any prime local dimension in the adversarial scenario, which can be applied immediately to verifying measurement-based quantum computation. Our protocol requires only local Pauli measurements and is easy to realize with current technologies. It achieves the optimal scaling behaviors with respect to the system size and the target precision, and exponentially enhances the scaling behavior with the significance level.
Authors: Zihao Li, Huangjun Zhu and Masahito Hayashi

Title: Scalable Non-Stabilizerness Recognition with Machine Learning
Abstract: Quantum computing's promise lies in leveraging unique quantum resources, with non-stabilizerness (or magic) being crucial for achieving computational power beyond classical capabilities. However, detecting and quantifying magic in large, entangled systems remains a significant challenge. In this paper, we present a machine learning framework based on Convolutional Neural Networks (CNNs), designed to efficiently classify stabilizer and non-stabilizer states using partial information acquired from measurement outcomes. By introducing a simple yet effective sorting step for measurement outcomes, we organize the data into a more structured format, enabling the CNN to extract relevant features with higher accuracy than previous methods. This approach achieves robust performance on simulated quantum systems with up to 54 qubits, greatly extending the scalability of stabilizer classification. Moreover, the computational efficiency and experimental practicality of our sorting-based method make it well suited for real-world quantum devices.

Authors: Sixuan Wu, Zanqiu Shen, Chenghong Zhu, Guangxi Li and Xin Wang

Title: ShuttleFormer: A Machine Learning Approach to Shuttle Scheduling in Trapped-Ion

Abstract: Trapped-ion systems are emerging as a leading quantum hardware platform for fault-tolerant quantum computation, owing to their high gate fidelity, all-to-all connectivity, and long coherence times. A representative architectural model within this platform is the trapped-ion linear tape (TILT), which supports scalable quantum computing through modular trap arrays. Executing quantum algorithms on such hardware requires quantum compilers to bridge the gap between high-level algorithms and low-level physical operations, primarily through shuttling-based qubit movement. However, efficiently scheduling shuttling operations remains a key challenge, particularly as circuit sizes and hardware complexity grow.

In this work, we propose a machine learning–based approach to shuttle scheduling that leverages structural patterns learned from quantum circuits. By incorporating an attention mechanism tailored to the constraints of TILT devices, we develop ShuttleFormer. Our method outperforms existing heuristic-based compilers, achieving an average fidelity improvement of 49.29 percent on well-known benchmark applications. It also reduces compilation time by an average of 51.14 percent, with gains reaching up to 88.82 percent in complex circuits. These results highlight a promising direction for advancing compiler support in large-scale trapped-ion quantum computing.

Authors: Xiyao Feng, Chenghong Zhu, Xian Wu, Jingbo Wang, Guangxi Li and Xin Wang

Title: Signed Designs for Learning Quantum State Properties with Applications

Abstract: Frame theory provides an elegant framework for reconstructing Hilbert spaces, with established applications in quantum computation for state tomography. However, pragmatic implementations face challenges in constructing dual frame elements economically and realizing efficient circuit structures. Here, we address these issues by generalizing tight measurement frames to their signed counterparts and constructing accessible signed tight measurement frames. We demonstrate that they preserve many key features of the positive case at the expense of modified variances while admitting novel and simple constructions. We provide one concrete realization: a weighted signed 2-design ensemble featuring linear (in the number of qubits) CNOT gates, logarithmical circuit depth, and constant measurement overhead. Notably, this construction, which is promising for measuring off-diagonal observables efficiently, offers a significant improvement over state-of-the-art tight measurement frames.

Authors: Yi-Hsin Lin, Scott Smart and Prineha Narang

Title: Supervised binary classification of small-scale digit images and weighted graphs with a trapped-ion quantum processor

Abstract: In this work, we present the results of benchmarking a quantum processor based on trapped 171Yb+ ions by performing quantum machine learning algorithms. Using a quantum-enhanced support vector machine algorithm with up to five qubits we perform a supervised binary classification on two types of datasets: small binary digit images and weighted graphs with a ring topology. For the first dataset, images are intentionally selected so that they could be classified with 100% accuracy. This allows us to specifically examine different types of quantum encodings of the digit dataset and study the impact of experimental noise. In the second dataset, graphs are divided into two categories based on the spectral structure of their Ising Hamiltonian models, which is related to the NP-hard problem. We applied a QAOA-inspired encoding scheme that uses n of entangling gates for embedding the Ising spectrum of 2^n size into

the probability amplitudes of an entangled state, exploiting the full dimensionality of n-qubit Hilbert space. The encoding structure allows for consideration of non-optimized and optimized versions of the corresponding quantum-enhanced support vector machine circuits, consisting of 2n and n entangling gates, respectively. For both problems, we study various levels of circuit optimization and found that, for all experiments conducted, we achieve classifiers with 100% accuracy on both training and testing datasets. This demonstrates that the quantum processor has the ability to correctly solve the basic classification task under consideration.
Authors: Anastasiia Nikolaeva, Evgeniy Kiktenko, Ilia Zalivako, Alexander Gircha, Alexander Borisenko, Ilya Semerikov, Aleksey Fedorov and Nikolay Kolachevsky

Title: Tailor Made Embeddings For Quantum Machine Learning

Abstract: We present an autoencoder strategy to solve the embeddings of real-world large-sized classical data sets like IMAGENET on quantum circuit models. We show preliminary results where we can embed IMAGENET in an 11-qubit circuit and, with noise currently, recover the original image. The tailor-made embedding created by the autoencoder helps recover the underperformance of quantum machine learning models versus classical models, as shown in Better than classical? The subtle art of benchmarking quantum machine learning models.
Authors: Aldo Lamarre and Dominik Safranek

Title: The curse of random quantum data and a cure from Pauli distribution

Abstract: Quantum machine learning (QML) emerges as a transformative frontier in quantum computing. While existing studies focus on algorithm design, the role of quantum data is not yet fully understood. In this work, we establish a rigorous data-dependent framework for quantifying QML performance. We find that when the input quantum data is uniformly random, training and generalization capabilities will be exponentially suppressed by qubit numbers, which we term “the curse of random quantum data”. Furthermore, we demonstrate that structured, physically motivated distributions, for example, states generated by shallow quantum circuits exhibit biased Pauli coefficients that circumvent this limitation. This constitutes the first data-dependent training performance guarantee for QML, resolving a key gap in prior analyses that overlooked quantum state structure. Our results establish Pauli distribution analysis as a powerful framework for evaluating QML capacity, with immediate implications for analyzing the performance of quantum learning protocols with realistic and meaningful quantum states.
Authors: Kaining Zhang, Junyu Liu, Liu Liu, Liang Jiang, Min-Hsiu Hsieh and Dacheng Tao

Title: The Effectiveness of Classical and Hybrid Models for MaxCut problem

Abstract: The MaxCut problem is a fundamental NP-hard optimization task where the goal is to partition the vertices of an undirected graph into two disjoint subsets to maximize the total weight of the edges crossing the partition. Each solution corresponds to a binary string indicating vertex assignments, and its quality is measured by the approximation ratio—the ratio between the solution cost and the global optimum. MaxCut plays an important role in neural network applications, including clustering and partitioning in graph neural networks. It also relates to energy-based models such as Hopfield networks and Boltzmann machines, where low-energy configurations correspond to high-cut values. Moreover, methods originally developed for MaxCut, such as relaxation and approximation techniques, have influenced strategies for pruning and optimizing neural networks. Neural networks have recently shown strong capabilities in solving MaxCut directly by learning from data. These include supervised heuristics, reinforcement

learning, and unsupervised models like autoencoders and generative adversarial networks, which leverage structural properties of graphs to approximate good solutions efficiently. Quantum computing offers an alternative through variational quantum algorithms, with the Quantum Approximate Optimization Algorithm (QAOA) as a key example. Several QAOA enhancements have been proposed, but current quantum hardware imposes significant limitations on depth and precision. In this work, we investigate whether quantum neural network (QNN) models can achieve approximation ratios competitive with classical neural heuristics and QAOA-based methods, while offering advantages in terms of resource requirements. We compare five QNN architectures based on variational circuits with six classical neural models. Classical approaches display variable performance, with some models overestimating cut values and others underperforming. Quantum models converge more consistently but often reach lower approximation ratios and involve greater computational overhead. Both quantum and classical methods exhibit consistent types of error. Baseline strategies like the Goemans–Williamson approximation and QAOA still provide the most reliable results. However, Goemans–Williamson offers a fixed approximation guarantee, and QAOA requires fault-tolerant hardware to handle difficult instances effectively. Our findings suggest that quantum model design remains a critical challenge for MaxCut. While QAOA provides practical value under the right conditions, hybridizing quantum techniques with machine learning heuristics must be approached with caution, due to the complexity of required pre-training and the sensitivity of quantum methods to resource constraints.
Authors: Leonardo Lavagna, Francesca De Falco and Massimo Panella

Title: The Lie Algebra of XY-mixer Topologies and Warm Starting QAOA for Constrained Optimization

Abstract: The XY-mixer has widespread utilization in modern quantum computing, including in variational quantum algorithms, such as Quantum Alternating Operator Ansatz (QAOA). The XY ansatz is particularly useful for solving Cardinality Constrained Optimization tasks, a large class of important NP-hard problems. First, we give explicit decompositions of the dynamical Lie algebras (DLAs) associated with a variety of XY-mixer topologies. When these DLAs admit simple Lie algebra decompositions, they are efficiently trainable. An example of this scenario is a ring XY-mixer with arbitrary RZ gates. Conversely, when we allow for all-to-all XY-mixers or include RZZ gates, the DLAs grow exponentially and are no longer efficiently trainable. We provide numerical simulations showcasing these concepts on Portfolio Optimization, Sparsest k-Subgraph, and Graph Partitioning problems. These problems correspond to exponentially-large DLAs and we are able to warm-start these optimizations by pre-training on polynomial-sized DLAs by restricting the gate generators. This results in improved convergence to high quality optima of the original task, providing dramatic performance benefits in terms of solution sampling and approximation ratio on optimization tasks for both shared angle and multi-angle QAOA.
Authors: Steven Kordonowy and Hannes Leipold

Title: Toward Quantum Machine Translation via Quantum Natural Language Processing

Abstract: The present study is an investigation of the feasibility of language translation using quantum natural language processing algorithms on noisy intermediate-scale quantum (NISQ) devices. Classical methods in natural language processing (NLP) struggle with handling large-scale computations required for complex language tasks, but quantum NLP on NISQ devices uses quantum aspects to efficiently process and analyze vast amounts of linguistic data, potentially for NLP applications. Here, we are trying to make a model paving the way for quantum neural machine translation, which could potentially offer advantages over classical methods, for example for

language translation. Based on Shannon entropy, we demonstrate the significant role of some appropriate angles of rotation gates in the performance of parametrized quantum circuits. In particular, we use these angles (parameters) as a means of communication between quantum circuits of different languages. In QNLP, we use both bag-of-words models and compositional structure models. This work constitutes the first proof-of-concept for the compositional structure in QNLP, where both the parameters and the structure of quantum circuits are important for the interpretability of the model. To achieve this, we focus on selecting sentences with similar structures. To achieve our objective, we adopt the encoder-decoder model of classical neural networks and implement the translation task using long short-term memory (LSTM). Our experiments involved 160 samples comprising English sentences and their Persian translations. We trained the models with different optimizers implementing stochastic gradient descent (SGD) as primary and subsequently incorporating two additional optimizers in conjunction with SGD. We achieved optimal results, with mean absolute error of 0.03, mean squared error of 0.002, and 0.016 loss, by training the best model, consisting of two LSTM layers and using the Adam optimizer. Our small dataset, though consisting of simple synonymous sentences with word-to-word mappings, points to the utility of Shannon entropy as a figure of merit in more complex machine translation models for intricate sentence structures.
Authors: Mina Abbaszadeh, Mariam Zomorodi, Mehrnoosh Sadrzadeh, Vahid Salari and Philip Kurian

Title: Towards Optimization-Free Adaptive Ansatze

Abstract: Due to the precision required in quantum chemistry applications, measurement costs are a roadblock to the practical implementation of VQE algorithms, whose optimization requires a high number of cost function evaluations. In the case of ADAPT-VQE, one of the most popular algorithms for solving quantum chemistry problems using quantum computers, multiple full optimizations are required, as an ansatz is grown by selecting operators from a pool based on their gradients, with parameters being re-optimized after each addition. In this work, we propose an optimization-free adaptive algorithm that makes use of energy and gradient measurements to construct the ansatz. Available data is used to build a quadratic model for the cost and establish a search direction along which a new parameter vector is found. Our protocol slashes the measurement costs of ADAPT-VQE by replacing each full optimization with fewer than 10 energy evaluations. As compared to a similarly cost-frugal protocol where only the last parameter is optimized, which exhibits slow convergence and fails to reach chemical accuracy for strongly correlated systems, our protocol converges much faster and succeeds in all test cases.
Authors: Mafalda Ramôa, Luis Santos, Nicholas Mayhall, Edwin Barnes and Sophia Economou

Title: Towards quantum extreme learning and reservoir computing on utility-scale digital quantum processor

Abstract: Quantum Reservoir Computing (QRC) extends the classical reservoir computing paradigm to quantum systems, leveraging their complex and high-dimensional Hilbert space for temporal data processing. While QRC offers a promising path for quantum machine learning, most existing works have been focused on analog implementations and simulations. Practical deployment has been limited by the emergence of concentration phenomena that limit scalability.

In this work, we address these limitations by proposing a scalable, hardware-compatible QRC architecture designed for state-of-the-art digital quantum processors. Our approach emphasizes both theoretical foundations and experimental feasibility, incorporating a hybrid design that combines quantum feature encoding and classical feedback mechanisms. These elements

contribute to maintain expressivity while mitigating the effects of observable concentration and shot noise. A key contribution is a practical hyperparameter tuning strategy that identifies optimal regimes balancing robustness and model capacity. Crucially, the learning in our framework is performed offline: the quantum system evolves and is measured according to fixed dynamics, while the readout weights are trained only after the measurement outcomes are collected.

We evaluate our architecture on benchmark tasks, demonstrating strong performance and generalization without task-specific tuning. Indeed, by monitoring key metrics such as observable variability and output expressivity, we identify a universal regime of optimality that offers strong performance across diverse tasks and system sizes. This regime reflects the careful trade-off between the richness of quantum dynamics and the distinguishability of measurement outcomes.

Our findings contribute to the growing body of work on scalable QML by offering a practical and theoretically grounded QRC framework. The proposed architecture not only addresses current hardware constraints but also provides a foundation for future exploration of hybrid quantum-classical models. Overall, our results highlight QRC as a promising candidate for near-term quantum machine learning applications, particularly in tasks involving temporal data.

Authors: Timothee Dao, Ege Yilmaz, Ibrahim Shehzad, Christophe Pere, Kumar Ghosh, Corey O'Meara, Giorgio Cortiana, Stefan Woerner and Francesco Tacchino

Title: Trainability of Parameterised Linear Combinations of Unitaries

Abstract: A principal concern in the optimisation of parameterised quantum circuits is the presence of barren plateaus, which present fundamental challenges to the scalability of variational algorithms and quantum machine learning models. Recent proposals for these models have increasingly used the linear combination of unitaries (LCU) procedure as a core component. In this work, we prove that an LCU of trainable parameterised circuits is still trainable. We do so by analytically deriving the expression for the variance of the expectation when applying the LCU to a set of parameterized circuits.

We support our conclusions with numerical results on linear combinations of fermionic Gaussian unitaries (matchgate circuits). Our work shows that sums of trainable paramaterised circuits are still trainable, and thus provides a method to construct new families of trainable circuits. We conclude by showing that there is a scope for quantum advantage in these trainable circuits, focusing on the case of free fermion systems.

Authors: Nikhil Khatri, Stefan Zohren and Gabriel Matos

Title: Transferable Equivariant Quantum Circuits for TSP: Generalization Bounds and Empirical Validation

Abstract: In this work, we addressed the challenge of generalization in quantum reinforcement learning (QRL) for combinatorial optimization, focusing on the Traveling Salesman Problem (TSP). Training quantum policies on large TSP instances is often infeasible, so existing QRL approaches are limited to small-scale problems. To mitigate this, we employed Equivariant Quantum Circuits (EQCs) that respect the permutation symmetry of the TSP graph.

This symmetry-aware ansatz enabled zero-shot transfer of trained parameters from n-city training instances to larger m-city problems. Building on recent theory showing that equivariant architectures avoid barren plateaus and generalize well, we derived novel generalization bounds for the transfer setting. Our analysis introduces a term quantifying the structural dissimilarity between n- and m-node TSPs, yielding an upper bound on performance loss under transfer. Empirically, we trained EQC-based policies on small n-city TSPs and evaluated them on larger instances, finding that they retained strong performance zero-shot and further improved with

fine-tuning—consistent with classical observations of positive transfer between scales.
These results demonstrate that embedding permutation symmetry into quantum models yields scalable QRL solutions for combinatorial tasks, highlighting the crucial role of equivariance in transferable quantum learning.
Authors: Monit Sharma and Hoong Chuin Lau

Title: Unified Framework for Matchgate Classical Shadows

Abstract: Estimating quantum fermionic properties is a computationally difficult yet crucial task for the study of electronic systems. Recent developments have begun to address this challenge by introducing classical shadows protocols relying on sampling of Fermionic Gaussian Unitaries (FGUs), a class of transformations in fermionic space which can be conveniently mapped to matchgate circuits. The different protocols proposed in the literature use different sub-ensembles of the orthogonal group O(2n) to which FGUs can be associated. We propose an approach that unifies these different protocols. We begin by demonstrating a novel result generalizing the 3-design property of the Clifford group to a class of unitary ensembles composed of products of independent random Pauli rotations. Building on this result, we then prove the equivalence of the previous protocols and derive an optimal sampling scheme for the associated FGU ensembles. Due to their generality, our results may prove useful for other quantum information tasks such as benchmarking and designing initialization strategies for variational quantum circuits.
Authors: Valentin Heyraud, Héloise Chomet and Jules Tilly

Title: Unifying non-Markovian characterisation with an efficient and self-consistent framework

Abstract: Noise on quantum devices is much more complex than it is commonly given credit. Far from usual models of decoherence, nearly all quantum devices are plagued both by a continuum of environments and temporal instabilities. These induce noisy quantum and classical correlations at the level of the circuit. The relevant spatiotemporal effects are difficult enough to understand, let alone combat. There is presently a lack of either scalable or complete methods to address the phenomena responsible for scrambling and loss of quantum information. Here, we make deep strides to remedy this problem. We establish a theoretical framework that uniformly incorporates and classifies all non-Markovian phenomena. Our framework is universal, assumes no parameter values, and is written entirely in terms of experimentally accessible circuit-level quantities. We formulate an efficient reconstruction using tensor network learning, allowing also for easy modularisation and simplification based on the expected physics of the system. This is then demonstrated through both extensive numerical studies and implementations on IBM Quantum devices, estimating a comprehensive set of spacetime correlations. Finally, we conclude our analysis with applications thereof to the efficacy of control techniques to counteract these effects, including noise-aware circuit compilation and optimised dynamical decoupling. We find significant improvements are possible in the diamond norm and average gate fidelity of arbitrary SU(4) operations, as well as related decoupling improvements in contrast to off-the-shelf schemes. This work is based on Phys. Rev. X 15, 021047 (2025).
Authors: Gregory White, Petar Jurcevic, Charles Hill and Kavan Modi

Title: VAE-QWGAN: Addressing Mode Collapse in Quantum GANs via Autoencoding Prior

Abstract: Recent proposals for quantum generative adversarial networks (GANs) suffer from the issue of mode collapse, analogous to classical GANs, wherein the distribution learnt by the GAN

fails to capture the high mode complexities of the target distribution. Mode collapse can arise due to the use of uninformed prior distributions in the generative learning task. To alleviate the issue of mode collapse for quantum GANs, this work presents a novel hybrid quantum-classical generative model, the VAE-QWGAN, which combines the strengths of a classical Variational AutoEncoder (VAE) with a hybrid Quantum Wasserstein GAN (QWGAN). The VAE-QWGAN fuses the VAE decoder and QWGAN generator into a single quantum model, and utilises the VAE encoder for data-dependent latent vector sampling during training. This in turn enhances the diversity and quality of generated images. To generate new data from the trained model at inference, we sample from a Gaussian mixture model (GMM) prior that is learnt on the latent vectors generated during training. We conduct extensive experiments for image generation QGANs on MNIST/Fashion-MNIST datasets and compute a range of metrics that measure the diversity and quality of generated samples. We show that VAE-QWGAN demonstrates significant improvement over existing QGAN approaches.
Authors: Aaron Thomas, Harry Youel and Sharu Jose

Title: Wavelet vision transformers and quantum pyramidal networks for biomedical image analysis

Abstract: We present an innovative hybrid quantum-classical architecture designed for biomedical imaging applications, specifically targeting pulmonary nodule classification in computed tomography (CT) images. Our approach integrates the advanced multi-scale feature extraction capabilities of the Wavelet Vision Transformer (Wave-ViT) with the computational strengths of a quantum orthogonal pyramidal circuit built from reconfigurable beam splitter (RBS) gates.

Initially, the pretrained Wave-ViT efficiently captures multi-scale, hierarchical features using wavelet-enhanced self-attention mechanisms, resulting in rich 128-dimensional embeddings representative of critical diagnostic features. We apply PCA to significantly reduce these embeddings' dimensionality, preparing them for efficient quantum processing.

The reduced feature vectors then serve as input to our quantum circuit, which leverages orthogonal transformations via RBS gates. This quantum pyramidal network efficiently processes information even with few qubits.

Evaluations conducted using the well-established LIDC-IDRI lung nodule dataset reveal that our quantum hybrid method matches or surpasses the accuracy, F1-score, and area under the ROC curve (AUC-ROC) of the classical Wave-ViT model. Notably, our approach excels particularly in more challenging imaging planes, highlighting improved robustness and feature representation. The successful reduction of dimensionality and complexity, combined with maintained or enhanced diagnostic performance, underscores the potential of quantum-classical hybrid systems to revolutionize medical imaging analysis by balancing computational efficiency with diagnostic accuracy.
Authors: Xavi Font Aragones and Miguel Ángel González Ballester

Title: When Quantum and Classical Models Disagree: Learning Beyond Minimum Norm Least Square

Abstract: Quantum Machine Learning Algorithms based on Variational Quantum Circuits (VQCs) are important candidates for useful application of quantum computing. It is known that a VQC is a linear model in a feature space determined by its architecture. Such models can be compared to classical ones using various sets of tools, and surrogate models designed to classically approximate their results were proposed. At the same time, quantum advantages for learning tasks have been proven in the case of discrete data distributions and cryptography primitives. In this work, we

propose a general theory of quantum advantages for regression problems. Using previous results, we establish conditions on the weight vectors of the quantum models that are necessary to avoid dequantization. We show that this theory is compatible with previously proven quantum advantages on discrete inputs, and provides examples of advantages for continuous inputs. This separation is connected to large weight vector norm, and we suggest that this can only happen with a high dimensional feature map. Our results demonstrate that it is possible to design quantum models that cannot be classically approximated with good generalization. Finally, we discuss how concentration issues must be considered to design such instances. We expect that our work will be a starting point to design near-term quantum models that avoid dequantization methods by ensuring non-classical convergence properties, and to identify existing quantum models that can be classically approximated.
Authors: Slimane Thabet, Léo Monbroussou, Eliott Z. Mamon and Jonas Landman

Title: A distributed approach to quantum approximate optimization

Abstract: The quantum approximate optimization algorithm (QAOA) is a variational quantum algorithm designed to obtain proximate solutions to combinatorial optimization problems. Despite great promise, the number of qubits it requires to solve large problem instances is far greater than what is currently available on any quantum computer. This led to the development of distributed approaches that make use of multiple quantum computers that collaboratively can solve a single large instance. Additionally, these techniques offer the possibility of using a single small quantum computer to solve large instances. Unfortunately these approaches have limitations of their own: for example, some of the existing approaches are limited to Z2 symmetric problem Hamiltonians. In this work we introduce a distributed form of QAOA that accepts any diagonal Hamiltonian, and provides on average better quality solutions compared to the Goemans-Williamson algorithm, although with a yet to be determined computational overhead.
Authors: Morounfoluwa Obidare, Ernesto Campos and Daniil Rabinovich

Title: A Hybrid Quantum-Classical AI Approach for Scalable and Precise Prediction of Cell-Level Molecular Biomarkers from Histology

Abstract: Hematoxylin and eosin (H&E)-stained images are the gold standard collected in routine clinical settings for histopathological assessment, playing a critical role in diagnostic decisions that influence patient treatment and outcomes. Additionally, immunohistochemistry (IHC) is used in clinical settings to quantify one or two protein markers at single-cell level. The expression of protein markers distinguishes functionally distinct cell subtypes, such as immune cells that suppress or promote tumor growth, and is important in guiding treatment selection and predicting patient prognosis. While advanced multiplex techniques like COMET enable simultaneous detection of multiple protein markers on the same tissue slide for a comprehensive understanding of cell functionalities, they are costly and require specialized equipment. In contrast, H&E-stained slides are widely available and routinely used in clinical workflows, making them great alternatives for inferring cell subtypes at scale. Deep learning models have shown promise in H&E image analysis tasks including tumor grading and spatial arrangement analysis. However, most existing methods work at the tissue or patch level rather than cell level, which could provide more granular insights into the tumor microenvironment. Our lab has been actively exploring cell-level prediction on H&E images in prior works. One initiative, AI4HE-Spatial, uses H&E images to predict cell-level molecular markers with deep learning models, to enable large-scale spatial biomarker profiling from widely available H&E images. However, classical deep learning models often require largely annotated datasets and generalize poorly in high-variance and low-sample settings, which is a common challenge in medical imaging. To address this, hybrid quantum-classical models have been proposed, demonstrating performance on par with classical models in medical imaging tasks. This suggests that quantum circuits may enhance feature representations by leveraging quantum principles such as superposition and entanglement. In this work, we explore the task of classifying breast cancer cells into two subtypes, estrogen receptor (ER)+ and ER-. The dataset comprises single-cell H&E images of breast cancer tissue. Cells that express both ER and epithelial markers (CK/EpCAM) were classified as ER+, while all other cells were labelled as ER-. ER+ breast cancers are generally less aggressive and typically respond well to hormone therapy. We compared two approaches for ER status prediction. The hybrid quantum-classical models use quanvolutional neural networks, where small image patches are encoded into quantum states and processed through a quantum circuit. The quantum features from the quantum measurements are then used to train classical convolutional neural network (CNN) models for downstream classification. The classical deep learning approach uses classical inputs (image pixel values) for the CNN models. Across both CNN architectures tested (ResNet18 and VGG19), models trained using the quanvolutional approach achieved comparable performance to those trained on the classical inputs. ResNet18 trained on classical inputs yielded an F1 score of 0.905 ± 0.016, while quanvolutional ResNet18 achieved 0.895 ± 0.008. Although classical inputs currently lead in performance, using quantum features remains competitive and promising. For future work, we will explore Hybrid Quantum-Classical Neural Networks (H-QNNs) where trainable quantum circuits are embedded within classical models such as ResNet to enhance representation capacity.

Authors: Jiasheng Isaac Cheong, Marcia Zhang, Wei Kit Tan, Solomonraj Wilson, Tianqi Chen and Mai Chan Lau

Title: A Resource-Efficient Quantum Kernel for High-Dimensional Learning on NISQ Devices

Abstract: We introduce CPMap, a novel quantum feature map designed for efficient and scalable kernel-based learning on near-term quantum devices. CPMap enables the encoding of high-dimensional classical data into compact quantum circuits with significantly reduced gate requirements. Unlike conventional approaches such as the ZZFeatureMap, CPMap requires quadratically fewer CNOT gates.

We conduct extensive benchmarking across diverse datasets to evaluate CPMap's performance, demonstrating that it consistently matches or outperforms the best available quantum kernels, and even exceeds classical kernels in several classification tasks. We further validate CPMap on IBM's ibm_quebec and ibm_torino devices, where it retains strong discriminative power under realistic noise. Our results, including statistical significance testing, underscore CPMap's robustness and practical viability for quantum machine learning on today's hardware.

Authors: Utkarsh Singh, Marco Armenta, Jean-Frederic Laprade, Aaron Z. Goldberg and Khabat Heshami

Title: Active learning with quantics tensor networks

Abstract: Tensor networks are a computational technique originating in condensed matter physics, however in recent years they have been increasingly applied to machine learning, in particular in the context of more efficient classical models and model compression, as well as a natural language for quantum machine learning. Here we present a novel active learning algorithm for constructing models of black-box functions from very few samples, which is based on a DMRG-like local optimization of tensor networks, but utilizing the tensor cross-interpolation which bypasses the need for defining an energy. We underscore the potential of this method for scientific

computing, particularly in the context of fast differentiation/integration and Fourier transforms. We illustrate these features using examples of simple mathematical functions.
Authors: Krzysztof Bieniasz and Hans-Martin Rieser

Title: Addressing the Current Challenges of Quantum Machine Learning through Multi-Chip Ensembles
Abstract: Practical Quantum Machine Learning (QML) is challenged by noise, limited scalability, and poor trainability in Variational Quantum Circuits (VQCs) on current hardware. We propose a multi-chip ensemble VQC framework that systematically overcomes these hurdles. By partitioning high-dimensional computations across ensembles of smaller, independently operating quantum chips and leveraging controlled inter-chip entanglement boundaries, our approach demonstrably mitigates barren plateaus, enhances generalization, and uniquely reduces both quantum error bias and variance simultaneously without additional mitigation overhead. This allows for robust processing of large-scale data, as validated on standard benchmarks (MNIST, FashionMNIST, CIFAR-10) and a real-world PhysioNet EEG dataset, aligning with emerging modular quantum hardware and paving the way for more scalable QML.
Authors: Junghoon Justin Park, Jiook Cha, Samuel Yen-Chi Chen, Huan-Hsin Tseng and Shinjae Yoo

Title: Advanced for-loop for QML algorithm search
Abstract: In this work, we leverage Large Language Model-based Multi-Agent Systems (LLMMA) for automated search and optimization of Quantum Machine Learning (QML) algorithms. Inspired by Google DeepMind's FunSearch, the proposed system works on abstract level to iteratively generate and refine quantum transformations of classical machine learning algorithms (concepts), such as the Multi-Layer Perceptron, forward-forward and backpropagation algorithms. As a proof of concept, this work highlights the potential of agentic frameworks to systematically explore classical machine learning concepts and adapt them for quantum computing, paving the way for efficient and automated development of QML algorithms. Future directions include incorporating planning mechanisms and optimizing strategy in the search space for broader applications in quantum-enhanced machine learning.
Authors: Fute Wong

Title: Advances in Quantum Annealing: From PUBO Formulations to Practical Implementation
Abstract: Quantum annealing represents a promising class of algorithms for solving optimization problems on quantum computing hardware. Many industrially relevant optimization problems can be naturally formulated as quadratic unconstrained binary optimization (QUBO) problems, which translate directly into quantum cost Hamiltonians where the ground state encodes the optimal solution. The quantum annealing protocol begins with a known initial state defined by a driver Hamiltonian that introduces quantum fluctuations to the system. The system parameters then get gradually tuned to reach the desired ground state. This evolution can proceed adiabatically—maintaining the system in the instantaneous ground state throughout—or diabatically, exploiting transitions into higher excited states to enable faster annealing times. Furthermore, quantum annealing schemes can be implemented on gate-based quantum computers through discretized protocols, as exemplified by the Quantum Approximate Optimization Algorithm (QAOA).
This work explores multiple strategies for enhancing quantum annealing performance. We investigate incorporating higher-order interaction terms into cost Hamiltonian formulations,

extending standard QUBO problems to polynomial unconstrained binary optimization (PUBO). Our analysis demonstrates that many optimization problems naturally admit PUBO formulations, yielding significant resource savings and potentially exponential reductions in annealing time compared to standard QUBO approaches. We validate these advantages numerically using the paradigmatic 3-SAT problem and show that those advantages persist even on gate-based quantum computers, where higher-order interaction terms require decomposition into the respective native gate sets.
We also address practical implementation challenges on current quantum hardware. Through analysis of an industrially relevant QUBO problem, we identify key noise sources in contemporary experimental systems encountered during annealing protocols. We demonstrate effective noise mitigation strategies that can be applied throughout the annealing sweep to completely cancel unwanted noise terms, thereby recovering ideal quantum annealing performance. Our findings are supported by analytical calculations that reveal universal scaling laws, extending the applicability of our results to general annealing problems.
Collectively, our results establish a promising pathway toward enhanced resource efficiency, reduced annealing times, and effective noise mitigation for quantum annealing protocols across analog and digital platforms. These advances represent crucial steps toward solving larger-scale optimization problems with industrial relevance.
Authors: Sebastian Nagies, Chiara Capecci, Kevin T. Geier, Marcel Seelbach Benkner, Javed Akram, Sebastian Rubbert, Dimitrios Bantounas, Michael Moeller, Michael Johanning and Philipp Hauke

Title: Advantage of Quantum Machine Learning from General Computational Advantage
Abstract: An overarching milestone of quantum machine learning (QML) is to demonstrate the advantage of QML over all possible classical learning methods in accelerating a common type of learning task as represented by supervised learning from classical data. However, the provable advantages of QML in supervised learning have been known so far only for the learning tasks designed for using the advantage of specific quantum algorithms, i.e., Shor's algorithms. Here we explicitly construct an unprecedentedly broader family of supervised learning tasks with classical data to offer the provable advantage of QML based on general quantum computational advantages, progressing beyond Shor's algorithms. Our learning task is feasibly achievable by executing a general class of functions that can be computed efficiently in polynomial time for a large fraction of inputs by arbitrary quantum algorithms but not by any classical algorithm. We prove the hardness of achieving this learning task for any possible polynomial-time classical learning method. We also clarify protocols for preparing the classical data to demonstrate this learning task in experiments. These results open vast opportunities to exploit a variety of quantum advantages in computing functions for the realization of provably advantageous QML.
Authors: Hayata Yamasaki, Natsuto Isogai and Mio Murao

Title: Agnostic Process Tomography
Abstract: Characterizing a quantum system by learning its state or evolution is a fundamental problem in quantum physics and learning theory with a myriad of applications. Recently, as a new approach to this problem, the task of agnostic state tomography was defined, in which one aims to approximate an arbitrary quantum state by a simpler one in a given class. Generalizing this notion to quantum processes, we initiate the study of agnostic process tomography: given query access to an unknown quantum channel and a known concept class C of channels, output a quantum channel that approximates the unknown channel as well as any channel in C, up to some error. In this work, we propose several natural applications for this new task in quantum machine learning, quantum metrology, classical simulation, and error mitigation. In addition, we give efficient

agnostic process tomography algorithms for a wide variety of concept classes, including Pauli strings, Pauli channels, quantum junta channels, low-degree channels, and a class of channels produced by QAC0 circuits. The main technical tool we use is Pauli spectrum analysis of operators and superoperators. We also prove that, using ancilla qubits, any agnostic state tomography algorithm can be extended to one solving agnostic process tomography for a compatible concept class of unitaries, immediately giving us efficient agnostic learning algorithms for Clifford circuits, Clifford circuits with few T gates, and circuits consisting of a tensor product of single-qubit gates. Together, our results provide insight into the conditions and new algorithms necessary to extend the learnability of a concept class from the standard tomographic setting to the agnostic one. The full version of our submission is available on arXiv with identifier arxiv:2410.11957.
Authors: Chirag Wadhwa, Laura Lewis, Elham Kashefi and Mina Doosti

Title: Almost fault--tolerant quantum machine learning with drastic overhead reduction

Abstract: The utility, or even feasibility, of supervised quantum machine learning has been doubted in various aspects, including its classical simulability, barren plateaus, and noise. This work addresses the issue of noise-induced barren plateaus, in which provably trainable models for ordinary barren plateaus no longer work, while other mitigation protocols are self-adaptive strategies either ineffective against depolarising noise, or require large overhead for noise characterisation. To overcome this, we introduce the partial quantum error correction protocol, where the distillations are abandoned to account for the high spacetime overhead. Without two-qubit gate errors, we simulated the training process of the number classification task based on MNIST datasets using variational circuits, where a depolarising channel is added to every single-qubit trainable unitary. We show that the model can still be trained even at a noise level p = 1.47 × 10^-3, which corresponds to a gate error rate of 1.96 × 10^-3 under randomised benchmarking, a value higher than that of state-of-the-art quantum computers. We further quantitatively investigate the variation in the loss function caused by the statistical fluctuations of the expectation values and determine the upper bound of non-trainability. The utility of this protocol is reinforced by careful analysis of the space-time cost of distillation, as well as comparing the model performance to that when the two-qubit gate errors are activated.
Authors: Haiyue Kang, Younghun Kim, Eromanga Adermann, Martin Sevior and Muhammad Usman

Title: Analyzing Generalization Error in Quantum Kernel Methods using Random Matrix Theory

Abstract: Understanding generalization is one of the central challenges in machine learning. While the double descent phenomenon is increasingly well understood in classical machine learning, its manifestation in the quantum domain remains largely unexplored and has only just started to be studied. Recent theoretical work has taken the first steps, particularly for ridgeless quantum kernel methods (QKM), revealing a sharp double descent peak. However, this initial analysis does not consider the crucial role of explicit regularization in ridge QKM, a standard tool for controlling model complexity and improving performance in practical QKM implementations. This leaves a critical question unanswered: how does explicit regularization reshape the generalization landscape, and can we theoretically predict the double descent peak in ridge QKM?
In this work, we develop a rigorous theoretical framework to answer this question. By introducing a new formulation, we map the regularized QKM problem to an equivalent high-dimensional linear ridge regression problem. This mapping allows us to apply powerful tools from random matrix theory (RMT) to derive a closed-form, analytical expression for the generalization error.
Our primary result is a precise formula that quantitatively describes the generalization error across the entire parameterization spectrum as a function of the ratio of model parameters to training

samples, and crucially, the regularization coefficient. This formula explicitly demonstrates how the double descent peak at the interpolation threshold is progressively suppressed and smoothed as the regularization coefficient increases, providing a direct, tunable mechanism for mitigating overfitting.
These theoretical predictions are corroborated by numerical experiments, which agree with our analytical formulas, particularly in predicting the peak behavior at the interpolation threshold even for modest model sizes. The key impact of our work is to move the understanding of QKM generalization from a qualitative observation for a ridgeless case to a quantitative, predictive framework that incorporates explicit regularization. Beyond providing new theoretical insight into QKM behavior, our framework offers practical guidelines for selecting regularization hyperparameters without exhaustive grid searches. Ultimately, this work provides a deeper, more predictive understanding of how to build robust and reliable QKM models, paving the way for more effective applications of quantum machine learning.
Authors: Kensuke Kamisoyama, Lento Nagano and Koji Terashi

Title: Applications of Quantum Convolutional Neural Networks in Medical Image Processing

Abstract: Machine learning and artificial intelligence have become imperative in our modern society as seen by their wide range of applications in education, finance, manufacturing, and healthcare. One key area where machine learning has shown particular promise is through the use of convolutional neural networks (CNNs) in medical image analysis and recognition, since CNNs can capture local spatial correlations in input image data and extract a large number of features from it. In light of this, the advancements of quantum machine learning have introduced a variety of quantum ansatze for classical CNNs, called quantum convolutional neural networks (QCNNs), some of which are capable of capturing global correlations in input image data by leveraging entanglement. In this work, various QCNN models are trained, validated, and tested on the low-resolution BreastMNIST and PneumoniaMNIST grayscale image datasets, in a binary classification task. This work benchmarks different quantum data encoding strategies and a variety of variational quantum circuit architectures that can be used to construct the QCNN model. The circuits are built using the open-source Python package, HierarQcal, which is a quantum circuit builder that simplifies the process of quantum circuit design, composition, generation, scaling, and trainable parameter management. This package includes a range of circuit structure hyperparameters that govern the connectivity of the parameterised two-qubit unitary gates that are used in the convolution and pooling layers of the QCNN architecture. This work focuses on how the performance of a QCNN model changed depending on the choice of unitary gate and its corresponding number of trainable parameters. In addition to this, varying the circuit structure hyperparameters also indicated that the performance of a QCNN model is not only determined by the choice of the unitary gate and its number of parameters, but also by the connectivity of the gates throughout the circuit structure. These findings contribute to a more systematic understanding of QCNN design principles and demonstrate that careful tuning of both quantum operations and circuit connectivity is crucial for optimising learning in near-term quantum devices.
Authors: Dhiya Dharampal, Francesco Petruccione and Ilya Sinayskiy

Title: Average-Case Algorithms for Local Hamiltonian Problem

Abstract: We introduce quantum algorithms for optimizing local Hamiltonian problems defined on graphs. We develop formulae to analyze the average-case energy achieved by these algorithms on random regular graphs in the infinite-size limit. We compare these algorithms to simple classical approaches and a state-of-the-art worst-case algorithm. We find that our algorithms on average

outperform these approaches for the EPR Hamiltonian on random regular graphs, and the Quantum MaxCut (QMC) Hamiltonian on random regular bipartite graphs. As a special case, we show that our algorithms prepare states within 1.62% error of the ground state energy for QMC on an infinite 1D ring, corresponding to the quantum Heisenberg spin chain.
Authors: James Sud, Kunal Marwaha and Adrian She

Title: Barren-plateau free variational quantum simulation of Z2 lattice gauge theories
Abstract: In this work, we use a variational quantum eigensolver (VQE) to investigate ground states and static string breaking in a Z2 lattice gauge theory. We consider a two-leg ladder lattice coupled to Kogut-Susskind staggered fermions. Simulations using tensor networks are used to verify the VQE results. We find that for varying Hamiltonian parameter regimes, and in the presence of external charges, the VQE is able to arrive at the gauge-invariant ground state without explicitly enforcing gauge invariance through penalty terms. For the theory involving charges, the VQE performs well, while the tensor network approach arrives at non gauge-invariant local minima. Thus, VQE is seen to be a promising tool for Z2 LGTs, and could pave the way for studies of other gauge groups.
Authors: Fariha Azad, Matteo Inajetovic, Stefan Kühn and Anna Pappa

Title: Bayesian Quantum Amplitude Estimation
Abstract: Quantum amplitude estimation (QAE) is a fundamental routine that offers a quadratic speed-up over classical approaches. The original QAE protocol is based on phase estimation. The associated circuit depth and width, and the assumptions of fault tolerance, are unfavorable for near-term quantum technology. Subsequent approaches attempt to replace the original protocol with hybrid iterative quantum-classical strategies, relying on simpler quantum circuits with m non-controlled applications of the amplification operator. These circuits are inserted in a classical feedback loop where a CPU chooses m for each iteration.
In this work, we present BAE, a problem-tailored and noise-aware Bayesian algorithm for quantum amplitude estimation capable of saturating the Heisenberg limit in a fault tolerant scenario. If device noise is present, BAE can dynamically characterize it and adapt in real-time. Our algorithm is based on greedy Bayesian inference, with a heuristic optimization routine that adaptively defines the experimental search range. We further propose aBAE, an annealed variant of BAE drawing on methods from statistical inference, to enhance robustness. Our proposals are parallelizable in both quantum and classical components, offer tools for fast noise model assessment, and can leverage preexisting information. Additionally, they accommodate experimental limitations and preferred cost trade-offs.
We propose a robust problem-agnostic benchmark for amplitude estimation algorithms and use it to test BAE against other QAE algorithms, demonstrating its competitive performance in both noisy and noiseless scenarios. In both cases, as compared to other algorithms, it achieves lower error for any cost, and lower cost for any error. Additionally, in the presence of decoherence, it is capable of learning when other algorithms fail or become erratic due to noise.
Even though it can adapt to noisy devices, BAE is not exclusively a NISQ algorithm: it is capable of achieving full quantum advantage in a fault-tolerant scenario. This makes it especially interesting in the transition between the NISQ and fault tolerant eras, as an algorithm that can interpolate between those regimes. In particular, BAE is capable of characterizing noise and self-adapting accordingly, seeking the best results given the limitations of the quantum hardware. While noise can still slow down the learning rate, proper handling can minimize this slowdown while safeguarding correctness. Our algorithm continues to learn even after others stagnate due to

noise, still displaying quantum-enhanced estimation.
Authors: Alexandra Ramôa and Luis Santos

Title: Characterization of distillable bipartite system enhanced by collective measurement and machine learning

Abstract: Entanglement distillation is a fundamental process in quantum information theory, allowing for the extraction of maximally entangled states from multiple copies of a given quantum state. We move forward the vastly studied 2×2 and 2×3 as higher dimensional systems present a richer geometry allowing better distillation of maximally entangled states. While sufficient conditions for the distillability have been studied, they often require the knowledge of the density matrix and so they are impractical as an experimental tool. Here, we explore the use of machine learning for the classification of the first non-trivial distillable system for dimensions $2 \times N$ (with $N > 3$). This paves the road for the exploration of systems of higher dimensions, which is an almost unexplored territory.
Authors: Antonio Mandarino, Christian Candeago, Paolo Da Rold, Michele Grossi and Pawel Horodecki

Title: Circuit compression for 2D quantum dynamics

Abstract: As quantum processors progress toward practical utility, the need for efficient compilation algorithms grows. In particular, simulating real-time dynamics of 2D systems with limited quantum resources requires the compression of accurate, deep circuits into shallow, hardware-friendly forms that retain high levels of accuracy. We present a variational approach for compiling deep time evolution operators into optimized shallower ansaetze using the Pauli propagation framework. Our method minimizes a fidelity-based cost function over special classes of ensembles of product states, which relate to the Hilbert-Schmidt cost by a well-known equivalence bound. By avoiding reliance on global Haar randomness and leveraging locally scrambled distributions, we achieve scalable circuit compression suited for near-term architectures, especially for systems with 2D connectivities.
Authors: Matteo D'Anna, Yuxuan Zhang, Roeland Wiersema and Juan Carrasquilla Alvarez

Title: Classical and Quantum Heuristics for the Binary Paint Shop Problem

Abstract: The Binary Paint Shop Problem (BPSP), a well-known APX-hard optimisation problem with significant applications in the automotive industry, has recently seen progress through the Quantum Approximate Optimisation Algorithm (QAOA). In particular, QAOA at depth $p = 7$ (QAOA7) has been shown to outperform classical heuristics on moderate-size instances. In this work, we explore two state-of-the-art QAOA variants—eXpressive QAOA (XQAOA) and Recursive QAOA (RQAOA)—and benchmark their performance on BPSP instances up to 4096 qubits. Our results show that XQAOA1 consistently outperforms RQAOA1 and classical heuristics across almost all instances. Notably, while RQAOA1's performance degrades with increasing problem size, approaching that of the best-known classical heuristic on the largest tested instance, XQAOA1 maintains robust performance and shows promise in surpassing all known heuristics in the asymptotic limit.
Authors: V Vijendran, Dax Enshan Koh, Ping Koy Lam and Syed M Assad

Title: Compilable QSVM with HHL in Qrisp

Abstract: Quantum support vector machines leveraging quantum linear systems algorithms like HHL have been proposed, providing an exponential speed-up for certain machine learning applications. To date, only a few implementations using manually constructed circuits exist, presumably due to significant challenges from a software engineering perspective regarding the proper coordination of quantum algorithmic primitives and management of quantum resources, and the need for real-time hybrid computations. In this work, we showcase a general end-to-end compilable implementation of a QSVM in the Qrisp programming framework based on our previous HHL implementation.

Authors: Matteo Inajetovic, Johannes Jung, Matic Petri?, Raphael Seidel and René Zander

Title: Convergence and Generalization of Warm-Starting Variational Quantum Algorithm

Abstract: We investigate the convergence and generalization of variational quantum algorithms (VQA) when an approximation of the global minimum (warm start) is available. This research is motivated by recent successes in guided local Hamiltonian problems, which convert a traditionally QMA-complete problem into a BQP-tractable one. While prior work has primarily focused on the trainability of warm-started variational quantum algorithms, this study examines their learnability, specifically analyzing how warm-starting influences the convergence and generalization properties of VQAs. Focusing on the task of predicting the ground state properties of a quantum many-body system, we start by proving that at the limit of infinite system size, the learning dynamics of the warm-starting variational algorithm with a constant-depth geometrically local ansatz are governed by a linear model obtained from the first-order Taylor expansion of the model around its initialization. While this theoretical result is exact in the infinite system size limit, we could approximate the outcomes of the algorithm from the linear model's result, even for finite, practical system sizes, due to the effect of warm starts. Such property allows us to access the model's convergence and generalization through its linearized version. We found that the warm start plays a critical role in controlling the generalizability and convergence efficiency of the model. Finally, we validate our findings with numerical experiments on 2D random Heisenberg models.

Authors: Tuyen Nguyen and Maria Kieferova

Title: Data Embedding on Two-Qubit interaction using Ising XYZ Hamiltonian model on Multiple Basis

Abstract: Noisy intermediate-scale quantum (NISQ) devices have limitations that hinder quantum machine learning algorithms, often leading to models being theoretical, tested on small toy datasets with noisy qubits, or simulated on classical devices. Consequently, rendering models useless for large datasets due to limitations on total feature values. We propose a data embedding method that allows for up to 3*k(k-1)/2 features for k qubits, while maintaining the expressivity of individual features. We embed features using the Ising interaction XYZ-based model Hamiltonian, with all-to-all connected qubits. Preliminary results highlight this method's ability to orthogonalize classes for binary classification when tested with Fidelity and Helstrøm Classifier methods on specific datasets. We present two use cases (image embedding and cluster pooling) using this embedding method, showcasing its usefulness for both classical and quantum machine learning.

Authors: Aakash Ravindra Shinde, Arianne Meijer-van de Griend, Ilmo Salmenperä, Valter Uotila and Jukka K. Nurminen

Title: Differentiable Digital Twins & Machine Learning for Quantum Computing

Abstract: Current techniques for characterisation of prototype quantum devices to gain insights into their performance and limitations are very manual, involving a lot of trial and error. This is a time-consuming endeavour that can take weeks or months to complete and often results in a sub-optimal understanding of true device performance. This is a major bottleneck to the development of quantum devices and is one of the key reasons why quantum technology is still in its infancy. In this paper, we present a novel approach to quantum device characterisation that uses machine learning and differentiable digital twins to automate the process. We demonstrate how this approach can be used to significantly speed up the characterisation process and how it can provide insights into the performance of quantum devices that were previously not possible.

Authors: Anurag Saha Roy

Title: DisCoCLIP: A Distributional Compositional Tensor Network Encoder for Vision-Language Understanding

Abstract: Vision-language understanding remains challenging, as models like CLIP often overlook linguistic structure and word order. We present DisCoCLIP, which combines a frozen CLIP image encoder with a text encoder that produces structure-aware sentence embeddings using tensor networks with linguistically informed architectures. Sentences are modeled as networks of word and composition tensors, with high-rank components factorized via Matrix Product States to control parameter growth. Trained using a self-supervised contrastive loss, DisCoCLIP matches or exceeds CLIP’s performance on benchmarks requiring sensitivity to verb usage and word order—achieving 82.42% on the SVO-Probes Verb subset (CLIP: 77.60%), 93.68% on SVO-Swap (CLIP: 57.89%), 67.58% on ARO Attribution (CLIP: 61.00%), and 55.12% on ARO Relation (CLIP: 51.53%), while using only 0.5 million trainable parameters compared to CLIP’s 63 million.

Authors: Kin Ian Lo, Hala Hawashin, Mina Abbazadeh, Tilen Limback-Stokin and Mehrnoosh Mehrnoosh Sadrzadeh

Title: Distilling Quantum Adversarial Manipulations via Classical Autoencoders

Abstract: Quantum neural networks have been proven robust against classical adversarial attacks, but their vulnerability against quantum adversarial attacks is still a challenging problem. Our work introduces a new quantum-classical machine learning framework and demonstrates its robustness against data manipulations which otherwise could easily trick quantum ML models. The implemented hybrid technique leverages the denoising capability of classical autoencoders to address an important challenge for the adaptation of quantum ML in security-sensitive applications such as self-driving vehicles and military systems. Our technique recovers quantum classifier accuracies when tested under standard machine learning benchmarks using MNIST and FMNIST image datasets, and PGD and FGSM adversarial attack settings. Our work highlights a promising pathway to achieve fully robust quantum machine learning in both classical and quantum adversarial scenarios. The demonstrated seamless integration of quantum and classical systems will open a new avenue for future research in which hybrid ML architectures can be designed leveraging key properties of both regimes.

Authors: Amena Khatun and Muhammad Usman

Title: Dynamic Estimation Loss Control in Variational Quantum Sensing via Online Conformal Inference

Abstract: Quantum sensing exploits non-classical effects to overcome limitations of classical sensors, with applications ranging from gravitational-wave detection to nanoscale imaging. However, practical quantum sensors built on noisy intermediate-scale quantum (NISQ) devices face significant noise and sampling constraints, and current variational quantum sensing (VQS) methods lack rigorous performance guarantees. This paper proposes an online control framework for VQS that dynamically updates the sensor's variational parameters while providing deterministic error bars on the estimates. By leveraging online conformal inference techniques, the approach produces sequential estimation sets with a guaranteed long-term risk level. Experiments on a quantum magnetometry task confirm that the proposed dynamic VQS approach maintains the required reliability (e.g., 90% coverage) over time, while still yielding precise estimates. The results demonstrate the practical benefits of combining variational quantum algorithms with online conformal inference to achieve reliable quantum sensing on NISQ devices.

Authors: Ivana Nikoloska, Hamdi Joudeh, Ruud van Sloun and Osvaldo Simeone

Title: Effect of Hybrid Model Structure on Reinforcement Learning Performance

Abstract: Hybrid quantum-classical neural networks have emerged as promising tools within reinforcement learning (RL), potentially offering enhanced performance and reduced model complexity compared to classical neural networks alone. However, the influence of variational quantum circuit (VQC) structure, particularly with regard to entanglement topology, remains poorly understood. In this work, we explore how different entanglement configurations (none, linear, cyclic, and fully entangled) affect learning in hybrid quantum-classical models applied to the classic CartPole environment. Using a mean reward of 160 as a benchmark for successful learning in the CartPole environment, our hybrid model achieved this performance with 50 parameters, compared to the 86 required by the classical model. Our focus is on settings where the quantum component contains more trainable parameters than the classical counterpart, allowing us to assess performance under a quantum dominant hybrid model. We find that linear and cyclic entanglement structures consistently enable stable learning, while fully entangled and unentangled circuits did not consistently achieve stable learning. These results suggest that quantum reinforcement learning is sensitive to circuit structure, and they provide practical guidance for the design of effective hybrid models.

Authors: Saad Amir, Anton Dekusar and Biswajit Basu

Title: Efficient Estimation of the Quantum Fisher Information Matrix in Commuting-Block Variational Circuits

Abstract: The Quantum Fisher Information Matrix (QFIM) is a fundamental quantity in various subfields of quantum physics. It plays a crucial role in the study of parameterized quantum states, as it quantifies their sensitivity to variations in their parameters. Recently, the QFIM has been successfully employed to enhance the optimization of variational quantum algorithms. However, its practical applicability is often hindered by the high resource requirements for its estimation.
In this work, we introduce a novel protocol for computing the off-block-diagonal elements of the QFIM between different layers in a particular class of variational quantum circuits, known as commuting-block circuits. Our approach significantly reduces the quantum resources required, specifically lowering the number of distinct quantum state preparations from $O(m^2)$ to $O(L^2)$, where m is the total number of parameters and L is the number of layers in the circuit. Consequently, our protocol also minimizes the number of classical measurements and

post-processing operations needed to estimate the QFIM, leading to a substantial improvement in computational efficiency.
Authors: Rafael Gomez Lurbe

Title: Efficient Offline Reinforcement Learning via Quantum Reward Encoding

Abstract: Efficiently and effectively utilizing offline reinforcement learning (RL) with limited samples has long been a challenging and critical issue. Learning effective policies with limited sample information aligns closely with real-world scenarios, yet the performance of offline RL frequently falls short. This study combines RL with quantum autoencoders to design a reward supervision learner under limited sample data. We construct a quantum reward encoder using a Parameterized Quantum Circuit (PQC), which enables data compression while performing reward supervision learning in the quantum state space. Our model is referred to as Quantum Reward Encoding (QRE), which efficiently learns quantum embeddings of states while simultaneously supervising the learning of rewards within a supervised framework. After training, we use QRE to obtain the quantum embedding of the state and decode the corresponding reward. We then replace the real states and rewards to facilitate the subsequent RL training. We conducted offline RL experiments on three well-known datasets with a limited sample size of 100 samples. We used Soft-Actor-Critic (SAC) and Implicit-Q-Learning (IQL) to demonstrate the effectiveness of our approach. We found that using quantum embeddings for states, with the decoded embeddings serving as rewards, significantly enhances RL performance. On average, we achieved a 115.9% improvement in maximum reward performance across the three datasets for SAC and 117.8% for IQL. Furthermore, we observed that the quantum-embedded states exhibit exceptionally low delta hyperbolicity. We believe this phenomenon contributes to the effectiveness of QRE. The low delta hyperbolicity and effectiveness of QRE provide valuable insights for developing efficient offline RL methods under limited sample conditions.
Authors: Yewei Yuan, Outongyi Lv and Nana Liu

Title: Efficient Quantum Convolutional Neural Networks for Image Classification: Overcoming Hardware Constraints

Abstract: While classical convolutional neural networks (CNNs) have revolutionized image classification, the emergence of quantum computing presents new opportunities for enhancing neural network architectures. Quantum CNNs (QCNNs) leverage quantum mechanical properties and hold potential to outperform classical approaches. However, their implementation on current noisy intermediate-scale quantum (NISQ) devices remains challenging due to hardware limitations. In our research, we address this challenge by introducing an encoding scheme that significantly reduces the input dimensionality, eliminating the need for classical dimensionality reduction pre-processing. We demonstrate that a primitive QCNN architecture with 49 qubits is sufficient to directly process 28×28 pixel MNIST images. Our approach demonstrates advantages in accuracy and convergence speed with a similar parameter count compared to optimized classical CNNs. We validated our experiments on IBM's Heron r2 quantum processor, achieving 96.08% classification accuracy, surpassing the 71.74% benchmark of traditional approaches under identical training conditions. To our knowledge, this is the first native full-resolution image classification on real quantum hardware, validating the potential of quantum computing in this area.
Authors: Peter Röseler, Oliver Schaudt, Helmut Berg, Christian Bauckhage and Matthias Koch

Title: Efficient State Preparation with Bucket Brigade QRAM and Segment Tree

Abstract: The preparation of data into quantum states is a critical part of the design of quantum algorithms. Efficient state preparation techniques are fundamental to avoid the data bottleneck, where the cost of preparing a quantum state is higher than the complexity of the algorithm itself. In particular, when developing quantum algorithms for data analysis it is necessary to load a matrix or vector onto a quantum register. To achieve this it is common to assume the existence of a classical data structure, known as KP-trees, and Quantum Random Access Memory (QRAM), the counterpart of a classical RAM. Despite being common assumptions, no prior work looks at either how to load the KP-trees onto a physical implementation of the QRAM, or develops an end-to-end algorithm to perform state preparation of matrices and vectors. In this work, we address this gap by first showing how a real-values matrix can be loaded onto a Segment Tree, a generalization of the KP-tree; then showing how this structure can be loaded on the Bucket Brigade QRAM (BBQRAM). The choice of the BBQRAM is motivated by two main factors: the BBQRAM is arranged as a binary tree, thus exhibiting a logarithmic depth in the number of indexed elements, and it has been shown to be more robust to errors and noise when compared to other architectures. Finally, we develop the algorithm to perform state preparation with the Segment Tree and the BBQRAM in polylogarithmic time, giving a detailed account of the circuit implementation. For future work, we intend to extend this framework to handle complex-valued matrices and explore pruning strategies for the Segment Tree to reduce BBQRAM size.

Authors: Francesco Ghisoni and Alessandro Berti

Title: Employing an Integrated MCDM with Quantum Fuzzy Logic in Next-Generation Communication

Abstract: The forthcoming 6G and next-generation networks necessitate dynamic multi-criteria decision support systems capable of learning, adapting, and elucidating under conditions of severe uncertainty. Quantum computing methods and fuzzy logic decision support systems enhance the optimization of communication networks in intricate technical contexts. This research introduces a flexible methodology that combines Quantum Fuzzy Logic (QFL) analysis with Multi-Criteria Decision Making (MCDM) methods. This study presents the QFL strategy, which employs a multi-criteria decision-making methodology to assess the efficacy and signal strength of base stations and cellular global identification location characteristics. A quantum fuzzy logic (QFL) approach is introduced, which encodes membership functions as qubit superpositions and assesses fuzzy rules via depth-optimized quantum circuits. The suggested approach utilizes the superposition and entanglement characteristics of quantum bits to assess large datasets and criteria with remarkable speed and precision. Our research seeks to prioritize security variables through the presented methodology, offering a novel perspective on hardware and software security evaluation in quantum computing. The study aims to identify and analyze the impact of barriers and to assess and rate alternative solutions to mitigate these barriers. The simulation (Qiskit + FL Python) will demonstrate the alteration and optimization of ratios in factors like efficiency, energy consumption, speed, and operation. These studies will demonstrate both standalone and integrated hybrid architectures through simulation-based evaluations, addressing the multi-objective issues in optimizing communication networks.

Authors: Emin Tarakci and Emine Can

Title: Energetic advantages for quantum agents in online execution of complex strategies

Abstract: Agents often execute complex strategies – adapting their response to each input stimulus depending on past observations and actions. Here, we derive the minimal energetic cost for

classical agents to execute a given strategy, highlighting that they must dissipate a certain amount of heat with each decision beyond Landauer's limit. We then prove that quantum agents can reduce this dissipation below classical limits. We establish the necessary and sufficient conditions for a strategy to guarantee quantum agents have energetic advantage, and illustrate settings where this advantage grows without bound. Our results establish a fundamental energetic advantage for agents utilizing quantum processing to enact complex adaptive behaviour.
Note: This submission is based on work in "Energetic advantages for quantum agents in online execution of complex strategies." arXiv:2503.19896 (2025), which has had preliminary acceptance into Physical Review Letters.
Authors: Jayne Thompson, Paul Riechers, Andrew Garner, Thomas Elliott and Mile Gu

Title: Enhancing Expressivity of Quantum Neural Networks Based on the SWAP test

Abstract: Quantum neural networks (QNNs) represent a promising class of hybrid quantum-classical architectures for machine learning applications. Current QNN approaches fall into two main categories: variational quantum circuits (VQCs) that often lack clear connections to classical machine learning models, and architectures that directly translate classical neural network components—particularly perceptrons—into quantum circuits. The latter approach could hold particular promise for leveraging the established success of classical neural networks into quantum computing applications.
We investigate a specific QNN architecture constructed exclusively from SWAP test circuits. Under amplitude encoding of classical inputs, this architecture becomes mathematically equivalent to a classical two-layer feedforward network with quadratic activation functions (Pastorello and Blanzieri, 2024). However, the original architecture suffers from fundamental expressivity limitations: it violates the universal approximation theorem due to its quadratic activation functions, constraining its learning capacity.
Through comprehensive numerical analysis across diverse real-world classical datasets, we demonstrate that despite these theoretical limitations, the architecture successfully learns many practical tasks. However, it fails on challenging synthetic benchmarks, specifically the parity check function and n-spiral task. We provide analytical proof that this failure stems from inherent limitations of the quadratic activation functions: The original two-layer QNN cannot learn parity check functions for input dimensions greater than two, regardless of network size.
To overcome these limitations, we introduce a modified QNN circuit utilizing a generalized SWAP test that incorporates multiple copies of input and weight quantum registers. All copies undergo the SWAP test with a single shared ancilla qubit, effectively creating an architecture analogous to classical neural networks with product layers. This circuit modification can alternatively be viewed as a standard SWAP test with generalized amplitude encoding.
Our enhanced architecture successfully learns parity check functions in arbitrary dimensions, solves the n-spiral task, and handles more challenging real-world datasets while maintaining the implementation simplicity that makes SWAP test circuits attractive for current quantum hardware. These results establish a framework for systematically enhancing QNN expressivity through classical task analysis and demonstrate that SWAP test-based architectures can achieve broad representational capacity, suggesting strong potential for both classical and quantum learning applications.
Several directions emerge for future investigation. First, we will explore alternative quantum encoding schemes beyond amplitude encoding to further enhance representational capacity while preserving circuit simplicity. Second, we plan to systematically address noise resilience, as decoherence effects represent a critical scalability challenge for these architectures. Finally, we will investigate Quantum Extreme Learning Machine (QELM) implementations based on our

generalized SWAP test circuit building blocks. By employing fixed random quantum weights and restricting training to classical output layers, QELMs could significantly reduce training complexity while maintaining competitive performance. This hybrid approach may enable quantum-parallel model execution, potentially offering computational advantages for large-scale quantum machine learning applications.
Authors: Sebastian Nagies, Emiliano Tolotti, Davide Pastorello and Enrico Blanzieri

Title: Error-mitigated quantum state tomography using neural networks
Abstract: State tomography is a widely used method for obtaining information about a state, but it is often susceptible to noise interference. In this work, we introduce a neural network-based tomography technique designed to mitigate the noise when detailed information about the noise is not available. We validate the effectiveness of our approach through various examples and analyze its characteristics. The results indicate that our algorithm is a valuable tool, capable of withstanding different types of noise and efficiently providing a relatively accurate estimate of the noise-free state.
Authors: Yixuan Hu, Mengru Ma and Jiangwei Shang

Title: Experimental data re-uploading with provable enhanced learning capabilities
Abstract: The last decades have seen the development of quantum machine learning, stemming from the intersection of quantum computing and machine learning. This field is particularly promising for the design of alternative quantum or quantum-inspired computation paradigms that could require fewer resources with respect to standard ones, for example in terms of energy consumption. In this context, we present the implementation of a data re-uploading scheme on a photonic integrated processor, applied to several image classification tasks, where it grants high accuracies. We thoroughly investigate the capabilities of this apparently simple model, which relies on the evolution of one-qubit states, by providing an analytical proof that our implementation is a universal classifier and an effective learner, capable of generalizing to new, unknown data. Hence, our results both demonstrate data re-uploading in a potentially resource-efficient optical implementation, as well as shed new theoretical insight into this algorithm, its trainability, and generalizability properties. This can pave the way to new, more resource-efficient machine learning algorithms, which might use our scheme as a subroutine.
Authors: Martin Mauser, Solène Four, Lena Marie Predl, Riccardo Albiero, Francesco Ceccarelli, Roberto Osellame, Philipp Petersen, Borivoje Dakic, Iris Agresti and Philip Walther

Title: Explaining Quanvolutional Neural Networks: A Frobenius Norm-Based Approach
Abstract: Neural networks for image classification are widely used in fields like security and healthcare, but their black-box nature makes explanations challenging. Existing explainability metrics fall short of providing an intuitive, task-specific aggregate measure. In this work, we introduce a Frobenius norm-based quantification for explainability and demonstrate its effectiveness on Quanvolutional Neural Networks (QuNNs), extending the study of explainability of hybrid quantum-classical learning models. Preliminary results suggest that the proposed metric offers valuable insights into QuNN's behavior, helping improve explainability measures.
Authors: Ritu Thombre and Lirandë Pira

Title: Fast quantum algorithms for PDEs with pseudospectral Fourier method via preconditioning

Abstract: We achieve polynomial speedups for solving Poisson and elliptic PDEs. We approximate respective continuous operators with their discretized versions using the pseudospectral Fourier differentiation method. Further, we decompose the operator as a product of an efficiently-invertible matrix and a sum of products of permutation matrices, enabling step-by-step inversion and bounding of the condition number with our novel condition number bounding techniques. Overall, our solvers yield polynomial improvement in dimension, and in particular, our Poisson solver achieves logarithmic dependence on precision compared to polylogarithmic dependence in previous work.

Authors: Mariia Sobchuk, Arsalan Motamedi, Grecia Castelazo and Pooya Ronagh

Title: Finding Lottery Tickets in Quantum Machine Learning: Sparse Trainability in Variational Quantum Circuits

Abstract: We investigate the lottery ticket hypothesis in the context of quantum machine learning. By pruning parameters in variational quantum circuits, we identify sparse subnetworks that retain essential structure while significantly reducing model complexity. Using learned parameter scores, we identify sparse circuits with significantly reduced parameter counts. Such compact architectures may enhance trainability, mitigate noise sensitivity, and improve the deployability of quantum models on near-term quantum devices.

Authors: Hyeondo Oh and Daniel Kyungdeock Park

Title: Full-Stack Assessment Framework for Quantum Machine Learning Models

Abstract: The current development of Quantum Machine Learning models follows a fragmented approach, with algorithmic design evaluated separately from implementation bottlenecks, overlooking critical factors like compilation overhead, hardware constraints, and energetics. We present a comprehensive full-stack assessment framework that systematically evaluates QML models across five critical dimensions: compilation overhead on diverse processor topologies, traffic characterization of quantum circuits, fidelity under realistic noise, expressibility of variational ansatze, and energy consumption on actual quantum hardware. Our methodology encompasses quantum kernel methods and quantum neural networks, compiled across several processor topologies with analytical fidelity modeling, expressibility analysis of ansatzes, and energetic considerations on IBM processors, critical for future cryo-CMOS systems with limited power budgets. Assessment through algorithm-level metrics alone provides an incomplete picture of model scalability and implementability, and this framework enables a multidimensional characterization of models across several implementation perspectives.

Authors: Rupayan Bhattacharjee, Sergi Abadal, Carmen G. Almudever and Eduard Alarcón

Title: Full-stack quantum machine learning on self-hosted quantum devices with Qiboml

Abstract: We present Qiboml, an open-source software library for Quantum Machine Learning (QML) integrated with the Qibo quantum computing framework. Qiboml interfaces with commonly used classical Machine Learning frameworks such as TensorFlow, PyTorch, and Jax. This combination enables users to construct quantum or hybrid classical-quantum models that can be executed on various hardware accelerators: multi-threading CPU, GPU, and multi-GPU for quantum simulation on classical hardware (using state-vector and tensor network approaches), as well as Quantum Processing Units (QPU) for execution on self-hosted quantum devices or remote devices

through different cloud providers. We showcase the capabilities of the library by presenting several useful QML applications.
Authors: Matteo Robbiati and Andrea Papaluca

Title: Generating Quantum Reservoir State Representations with Random Matrices

Abstract: We present a scalable framework for quantum reservoir computing (QRC) that leverages random measurement operators to extract rich, high-dimensional state representations from quantum devices. Motivated by the search for novel computing paradigms that exploit the intrinsic dynamics of physical systems, our work builds on the Extreme Learning Machine (ELM) philosophy, in which an untrained reservoir performs a fixed nonlinear embedding of input signals into a high-dimensional space, and only a low-dimensional linear readout is trained for the task. Quantum systems, owing to their exponential state-space scaling, are prime candidates for such reservoirs, with recent demonstrations exploring superconducting qubits, spin lattices, and trapped-atom arrays.
In traditional QRC implementations, the state representation of the reservoir is formed by measuring a limited set of observables on each qubit, often yielding relatively low-dimensional embeddings and limiting expressivity. Here, we replace this restricted measurement set with an ensemble of random Hermitian matrices drawn from well-studied Random Matrix Theory (RMT) ensembles. By projecting the reservoir's density matrix onto a large, fixed ensemble of random observables, we generate a much richer feature space. This approach is justified by the fact that, for reservoir computing, interpretability of individual features is secondary to their collective sensitivity to the reservoir's internal dynamics.
To illustrate the efficacy of our method, we study two prototypical quantum reservoirs: a five-atom Heisenberg spin chain driven by a time-dependent magnetic field, simulated via QuTiP, and a five-qubit gate-based reservoir employing angle encoding and inter-qubit entangling layers, simulated in PennyLane. In both systems, we generate state vectors by applying random Hermitian operators. We then train a linear readout on various benchmark tasks including time series prediction and interpolation.
Our results demonstrate that random-matrix measurements yield performance competitive with, and in some regimes superior to, standard Pauli-based readouts. We analyze the role of key reservoir parameters, such as coupling strength and measurement dimension, in determining predictive accuracy. In particular, we show that increasing the measurement dimension systematically improves performance up to a saturation point and reservoir tuning via coupling strength enables a trade-off between nonlinearity and memory depth.
Beyond numerical benchmarks, our work outlines clear paths toward experimental realization. By spatially decoupling measurement sites from input drives, as in many superconducting-qubit platforms, one can perform ensembles of non-destructive random measurements without perturbing the reservoir's primary dynamics. Ultimately, our random-matrix measurement protocol paves the way for compact, atomic-scale learning devices capable of real-world time-series processing.
Authors: Tobias Fellner, Samuel Tovey, Christian Holm and Michael Spannowsky

Title: Geometry-Aware Dictionary Learning and Quantum-Guided Bagging for Qubit-Efficient Recommender Systems

Abstract: Modern recommenders describe each item with hundreds of sparse semantic tags, yet most quantum pipelines still map one qubit per tag, demanding more than one hundred qubits—well beyond the reach of current noisy-intermediate-scale quantum (NISQ) devices and

prone to deep, error-amplifying circuits. We close this gap with a three-stage hybrid workflow. First, a geometry-aware low-rank and sparse dictionary learning compresses the tag space into clusters of atoms. Second, a shallow Quantum Approximate Optimisation Algorithm selects the five atoms that maximise a performance-driven QUBO combining DeltanDCG and DeltaAUC, using only a five-qubit register. Third, a 50-estimator bagged ensemble of shallow decision trees scores the resulting five-dimensional codes.
On the public QuantumCLEF 150-ICM benchmark, the five-qubit model achieves nDCG@10 = 0.4042, ROC-AUC = 0.8064, and Log-Loss = 0.0702, matching full-feature baselines while cutting qubit requirements by approximately 90 percent. These results show that NISQ hardware can be practically useful when quantum search is preceded by geometry-respecting compression and followed by noise-robust bagging.
Authors: Azadeh Alavi, Fatemeh Kouchmeshki, Abdolrahman Alavi, Jiayang Niu and Yongli Ren

Title: Graybox Approach for Qudit System Identification and Control

Abstract: Understanding and controlling engineered quantum systems is a vital step for realizing quantum technology applications. This is often a challenging task due to technological limitations, fabrication imperfections, and uncertainties such as environmental noise. Physics-based "Whitebox" methods that address those challenges usually work under some assumptions on the noise and/or the control. On the other hand, standard machine learning (blackbox) methods do not give any physical insights into the system. Here, we give a perspective on the "Graybox" approach that we have been developing recently; an approach that utilizes machine learning blackbox structures to model the physical quantities prone to uncertainties along with whitebox structures that encode quantum mechanical laws. We particularly focus on our recent results of designing a graybox structure for modelling a qudit system subject to a general unknown non-Markovian noise. Once trained, the graybox could be utilized in an optimization loop to find optimal control pulses for quantum gate implementation. Our method achieves high-fidelity gate optimisation across global and sublevel qudit gate classes, demonstrating the potential of the graybox approach in system identification and control.
Authors: Yule Mayevsky, Akram Youssry, Ritik Sareen, Gerardo Paz-Silva and Alberto Peruzzo

Title: Hybrid Learning and Optimization Methods for Solving Capacitated Vehicle Routing Problem

Abstract: The Capacitated Vehicle Routing Problem (CVRP) is a fundamental NP-hard problem in logistics. Augmented Lagrangian Methods (ALM) offer a flexible framework for solving constrained CVRP formulations, but their performance depends heavily on well-tuned penalty parameters. We propose a hybrid optimization approach that integrates deep reinforcement learning (RL) to automate penalty selection within both classical (RL-C-ALM) and quantum-enhanced (RL-Q-ALM) ALM solvers.
Using a Soft Actor-Critic agent, our method learns penalty values from CVRP instance features and constraint violations. In RL-Q-ALM, subproblems are encoded as QUBOs and solved using Variational Quantum Eigensolvers (VQE). The agent learns across episodes by maximizing solution feasibility and minimizing cost.
Experiments show that RL-C-ALM outperforms manually tuned ALM on synthetic and benchmark CVRP instances, achieving better solutions with fewer iterations. RL-Q-ALM matches classical solution quality on small instances but incurs higher runtimes due to quantum overhead. Our results highlight the potential of combining RL with classical and quantum solvers for scalable, adaptive combinatorial optimization.

Authors: Monit Sharma and Hoong Chuin Lau

Title: Hybrid Quantum-Classical Neural Networks with Data Reuploading for Binary Medical Image Classification

Abstract: We present a hybrid quantum-classical neural network architecture designed for binary classification of chest X-ray images from the MedMNIST dataset. Our architecture features a hybrid data reuploading scheme, wherein a classical convolutional encoder learns to preprocess image features, which are then reuploaded into a parameterized quantum circuit. This hybrid structure improves trainability and representation capacity compared to static encodings. We benchmark our approach on PneumoniaMNIST and demonstrate the feasibility of the scheme using limited qubit resources.

Authors: Jian Feng Kong, Chee Kwan Gan, Stefano Carrazza and Jun Yong Khoo

Title: Imbalanced classification with quantum kernel methods

Abstract: Support vector machines, as optimized binary classifiers, can be implemented on a quantum computer (QSVMs) with complexity logarithmic in the size of the vectors and the number of training examples. A QSVM is a nonsparse matrix exponentiation technique for efficiently performing a matrix inversion of the training data inner-product (kernel) matrix. However, due to hardware limitations on current noisy intermediate scale quantum (NISQ) devices, this approach poses significant bottlenecks. As a result, focus is shifting to quantum embeddings as a promising alternative for hybrid quantum-classical classification methods. Typical datasets from biomedicine are characterized by strong class imbalances and sparsity, posing challenges for conventional classifying models. By using a systematic approach to quantum embeddings, which define the model's ability to capture meaningful correlations in the data, the parameters for optimized classification results are investigated on synthetic datasets modeled to real-life datasets. The results show that the choice of the feature map can significantly affect performance, especially in identifying patterns within the minority class. We evaluate our approach on the publicly available and strongly imbalanced stroke prediction dataset, using two quantum embedding circuits. With a systematic reduction of the data size, we show that quantum kernel methods are able to perform at least as good as their classical counterparts in terms of expressivity without relying on oversampling techniques. We achieve AUC scores of 0.8 with the optimized quantum approach when using 10% of the initial dataset. The success of quantum embeddings in achieving competitive performance with reduced dataset sizes represents a validation of quantum machine learning algorithms in real-world applications. In addition, this approach addresses a fundamental machine learning challenge posed by the curse of dimensionality and the persistent requirement for extensive labeled training data. In summary, appropriately designed quantum embeddings for defined specific datasets enable QSVMs to match or even surpass classical SVM approaches.

Authors: Florian Heininger, Niels Halama, Sakshi Singh, Jeanette Miriam Lorenz and Matthias Weidemüller

Title: Improving Adaptive Variational Quantum Algorithms

Abstract: In this work, we discuss and compare leading proposals of ansätze for the electronic structure problem. Our main contributions are a new algorithm, termed CEO-ADAPT-VQE, and a comparison of key viability metrics for state-of-the-art ansätze: the CNOT count and measurement costs required to achieve a chemically accurate solution. In addition to our novel CEO-ADAPT-VQE algorithm, we consider several static VQEs: k-UpCCGSD, QNP, tUPS, oo-tUPS, and pp-tUPS.

From the comprehensive set of ansätze we simulated, ADAPT-VQE markedly appears as the most viable approach, with other ansätze either failing to reach chemical accuracy or requiring a prohibitive number of measurements—five orders of magnitude larger than required by ADAPT-VQE, including all costs incurred throughout the adaptive ansatz construction.
This surprising result contradicts the common belief that the adaptive ansatz construction incurs a measurement overhead: the optimization advantages enjoyed by ADAPT-VQE easily compensate for any additional measurements required by the adaptive protocol.
Authors: Mafalda Ramôa, Panagiotis Anastasiou, Luis Santos, Edwin Barnes, Nicholas Mayhall and Sophia Economou

Title: Improving Quantum Neural Networks performances by Noise-Induced Equalization

Abstract: Quantum noise poses significant limitations on the performance of current quantum computing systems. Even variational quantum machine learning models, designed to mitigate hardware noise, can be hampered by noise-induced barren plateaus that hinder optimal optimization. However, recent research suggests that carefully managed noise levels can positively impact generalization performance.
Here, we introduce a pre-training procedure to determine the quantum noise level that yields the most favorable landscape properties. At this level, noise induces what we define as an “equalization” of variational parameters: those that were previously the least influential gain more relevance, while the most influential ones are tempered. We characterize this phenomenon through the Quantum Fisher Information Matrix, deriving a practical recipe to estimate the noise strength that maximizes parameter equalization.
Our numerical simulations show that this optimal noise level is associated with improved generalization performance, suggesting a promising approach for enhancing the effectiveness of quantum machine learning algorithms.
Authors: Francesco Scala, Giacomo Guarnieri, Dario Gerace and Aurelien Lucchi

Title: Inductive Graph Representation Learning with Quantum Graph Neural Networks

Abstract: Quantum Graph Neural Networks (QGNNs) present a promising approach for combining quantum computing with graph-structured data processing. While classical Graph Neural Networks (GNNs) are renowned for their scalability and robustness, existing QGNNs often lack flexibility due to graph-specific quantum circuit designs, limiting their applicability to a narrower range of graph-structured problems, falling short of real-world scenarios. To address these limitations, we propose a versatile QGNN framework inspired by the classical GraphSAGE approach, utilizing quantum models as aggregators. In this work, we integrate established techniques for inductive representation learning on graphs with parametrized quantum convolutional and pooling layers, effectively bridging classical and quantum paradigms. The convolutional layer is flexible, enabling tailored designs for specific problems. Benchmarked on a node regression task with the QM9 dataset, we demonstrate that our framework successfully models a nontrivial molecular dataset, achieving performance comparable to classical GNNs. In particular, we show that our quantum approach exhibits robust generalization across molecules with varying numbers of atoms without requiring circuit modifications. In addition, the framework allows for the replacement of regular QGNN layers with attention layers in the form of Quantum Graph Attention Networks (QGATs), which are also benchmarked on QM9. Here, the attention layers are incorporated as trainable feature maps. Compared to our standard QGNNs, this approach not only facilitates the training on more samples with less computational resources, but also delivers significantly better results. Furthermore, we numerically investigate the scalability of the QGNN model. Specifically, we

numerically demonstrate the absence of barren plateaus in our quantum circuits as the number of qubits increases, suggesting that the proposed quantum framework can be extended to handle larger and more complex graph-based problems effectively.
Authors: Arthur Mendonça Faria, Ignacio Fernández Graña and Savvas Varsamopoulos

Title: Information-Minimal Two-Body Moment Learning for Scalable QUBO Optimisation

Abstract: Quadratic-unconstrained binary optimisation (QUBO) problems from logistics, finance, and machine learning often involve tens of thousands of binary variables—far exceeding today's NISQ-era hardware, which typically supports tens to hundreds of noisy qubits. Existing "qubit-efficient" quantum optimisation methods compress problem sizes but ultimately rely on heuristics that search for a single bit-string solution, incurring deep circuits, non-convex surrogates, or loose mean-field-like approximation guarantees.
We propose an orthogonal approach: we learn the relaxed distribution itself—its statistical shadow—directly modeling the first- and second-order statistical moments of the QUBO problem. Our circuit uses only $2 \log_2(N) + 2$ qubits (just 16 qubits at N=128) and captures the complete set of two-body marginals. Each measurement selects two indices via address registers and encodes their bits onto ancillas; a single vectorized pass (GPU-accelerable) aggregates shot statistics into the complete moment tensor (mu, nu), preserving full end-to-end differentiability.
To correct raw moments that violate consistency constraints, we softly project them toward the Sherali–Adams level-2 polytope—the tightest linear relaxation expressible from two-body data—using a differentiable KL iterative-proportional-fitting step. The repaired moments define a maximum-entropy Ising surrogate, from which hundreds of parallel Gibbs samplers efficiently produce bit-string solutions.
Results: On 300 random Erdős–Rényi Max-Cut instances (N=128), covering sparse to dense graphs (edge densities alpha in [0.25,128]), our best two solver variants maintain approximation ratios of at least 0.90 precisely in the regime where Gurobi's exact ILP solver hits its 600-second timeout, yet complete each instance in just ~87 seconds using a similar compute budget. Our method compresses $O(N^2)$ pairwise interactions into merely $O(\log^2(N))$ learned parameters and naturally generalizes to k-state Potts models, charting a principled, scalable path toward large-scale hybrid quantum–classical optimisation.
Outlook: Because our solver outputs the entire distribution over near-optimal bit-strings rather than a single heuristic solution, it implicitly constructs an exponentially compact embedding of the underlying graph structure—compressing the full N^2-variable problem into just $O(\log^2(N))$ parameters. This exponential informational bottleneck provides more than computational efficiency; it offers a fundamentally compact representation of graph properties. Such embeddings immediately open pathways to classical meta-learning, enabling rapid one-shot warm-start optimisation, instance-difficulty prediction, and deeper explorations of structural properties—exciting directions we leave open for future research.
Authors: Gordon Yuan Ning Ma, Ioannis Leonidas and Dimitris G. Angelakis

Title: Information-theoretic evaluation of quantum machine learning model complexity

Abstract: In quantum machine learning, the design of the quantum feature map is essential to achieve high learning performance. It motivates us to explore what and how certain quantum properties contribute to learning performance. It is however highly nontrivial to characterize the quantum feature map for this purpose. In this work, we propose a characterization method of the quantum feature map, inspired by an information-theoretic criterion for model selection in classical machine learning: the minimum description length principle. We compress unitary

matrices to evaluate their complexity and numerically show that the complexity captures the learning behaviour of a quantum machine learning model with different unitary maps.
Authors: Aoi Hayashi, Akitada Sakurai and Kae Nemoto

Title: Interference Beats Hallucination? A Controlled Study of Hybrid Quantum–Classical Language Models for Code Generation
Abstract: Large language models (LLMs) frequently "hallucinate" non-existent API calls when generating source code. We hypothesize that this failure mode is aggravated by the strictly linear inner transformations of classical transformers. Replacing even one linear component with a parameterised quantum circuit (PQC) introduces interference before Born-rule measurement, potentially suppressing self-inconsistent branches early. We build a minimal hybrid model in which the logits head of a 2-layer GPT-mini is swapped for a 6-qubit PQC and evaluate both models on 20 HumanEval-Lite coding tasks. Under identical training, the quantum head cuts the Invalid-Symbol Rate from 8.1% to 3.1% and shrinks the pre-collapse amplitude mass on invalid tokens (0.17 to 0.05; Wilcoxon $p < 2 \times 10^{-6}$). Compilation and edge-case metrics stay comparable, indicating that interference mainly prunes hallucinated identifiers rather than improving unseen logic. A single gradient step after a failing test fixes 40% of errors in the hybrid model versus 30% for the classical baseline. All code, data, and reproducibility artifacts accompany this paper.
Authors: Doron Podoleanu

Title: Interval-based Analysis of Variational Quantum Algorithms
Abstract: Variational Quantum Algorithms (VQA) are a class of hybrid quantum-classical algorithms based on parametrised quantum circuits. In these algorithms, classical data are first encoded into quantum states, which are then processed by a quantum circuit with free parameters. These parameters are tuned to optimize a scalar cost function for a given task, which is represented by the circuit expectation value. The training procedure is similar to that of classical neural networks, but without the activation functions, which VQAs are unable to implement due to their unitarity. VQAs inherit many of the challenges faced by classical deep learning models. Notably, they are vulnerable to adversarial inputs: small perturbations to the input quantum state may cause significant changes in the output prediction. In this work, we present an interval-based reachability analysis to formally verify the robustness of variational quantum circuits against adversarial perturbations. We show through an example how this approach, which is based on a technique originally developed for classical deep learning, can be applied in the quantum setting to address the problem of the formal verification of quantum circuits.
Authors: Nicola Assolini, Luca Marzari, Isabella Mastroeni and Alessandra Di Pierro

Title: Is the QAOA the Ultimate Solution for the MaxCut problem?
Abstract: The Quantum Approximate Optimization Algorithm (QAOA) is a variational quantum algorithm designed to address classically intractable combinatorial optimization problems. It alternates p layers of two unitaries: one derived from a problem-specific cost Hamiltonian and one based on a non-commuting mixer Hamiltonian. Starting from the ground state of the mixer, QAOA builds a quantum state parameterized by angle vectors, and optimizes the expectation value of the cost Hamiltonian using classical routines. Although effective in principle, QAOA suffers from scalability issues, noise sensitivity, and barren plateaus that reduce its performance on complex problem instances.
To overcome these limitations, we propose the Quantum Approximate Neural Solver (QANS), a

quantum neural network-based alternative inspired by variational methods. We apply QANS to the MaxCut problem, which involves partitioning the vertices of a graph into two subsets to maximize the number of edges between them—a well-known NP-hard task in quantum optimization. In our approach, an n-qubit register is initialized in the zero state and evolved through a parameterized, problem-agnostic quantum circuit without relying on a separate mixer Hamiltonian. As in QAOA, the goal is to maximize the expectation value of the cost function through parameter optimization. We validate QANS through experiments on various graph types including complete graphs, Erdős–Rényi graphs, complete binary trees, and random regular graphs, with up to 10 nodes. We test multiple QANS circuit structures, including rotation-only and universal ansatzes, all with a single layer, and compare them to QAOA implementations with up to 8 layers. The results show that QANS consistently achieves high approximation ratios—above 0.96 on average—even on larger graphs, outperforming vanilla QAOA by approximately 10 percent. Additional analysis of measurement outcomes confirms that QANS identifies optimal solutions more frequently, whereas QAOA tends to produce more suboptimal results. These findings highlight QANS as a promising alternative for solving MaxCut, combining low circuit depth with strong performance and resilience to barren plateaus.
Authors: Leonardo Lavagna, Francesca De Falco and Massimo Panella

Title: Language Model for Large-Text Transmission in Noisy Quantum Communications

Abstract: Quantum communication has the potential to revolutionize information processing, providing unparalleled security and increased capacity compared to its classical counterpart by using the principles of quantum mechanics. However, the presence of noise remains a major barrier to realizing these advantages. While strategies like quantum error correction and mitigation have been developed to address this challenge, they often come with substantial overhead in physical qubits or sample complexity, limiting their practicality for large-scale information transfer. Here, we present an alternative approach: applying machine learning frameworks from natural language processing to enhance the performance of noisy quantum communications, focusing on superdense coding. By employing bidirectional encoder representations from transformers (BERT), a model known for its capabilities in natural language processing, we demonstrate improvements in information transfer efficiency without resorting to conventional error correction or mitigation techniques. These results mark a step toward the practical realization of a scalable and resilient quantum internet.
Authors: Yuqi Li, Zhouhang Shi, Li Shen, Haitao Ma, Jinge Bao and Yunlong Xiao

Title: Learning junta distributions, quantum junta states, and QAC0 circuits

Abstract: In this work, we consider the problems of learning junta distributions, their quantum counterparts (quantum junta states), and QAC0 circuits, which we show to be close to juntas.
(1) Junta distributions. A probability distribution $p: \{-1, 1\}^n \rightarrow [0, 1]$ is a k-junta if it only depends on k bits. We show that they can be learned with error ε in total variation distance from $O(2^k \log(n)/\varepsilon^2)$ samples, which quadratically improves the upper bound of Aliakbarpour et al. (COLT'16) and matches their lower bound in every parameter.
(2) Junta states. We initiate the study of n-qubit states that are k-juntas, which are the tensor product of a k-qubit state and an $(n-k)$-qubit maximally mixed state. We show that these states can be learned with error ε in trace distance with $O(12^k \log(n)/\varepsilon^2)$ single copies. We also prove a lower bound of $\Omega((4k + \log(n))/\varepsilon^2)$ copies. Additionally, we show that, for constant k, $\Theta(2^n/\varepsilon^2)$ copies are necessary and sufficient to test whether a state is ε-close or 7ε-far from being a k-junta.
(3) QAC0 circuits. Nadimpalli et al. (STOC'24) recently showed that the Pauli spectrum of QAC0

circuits (with a limited number of auxiliary qubits) is concentrated on low degree. We remark that they implied something stronger, namely that the Choi states of those circuits are close to be juntas. As a consequence, we show that n-qubit QAC0 circuits with size s, depth d, and a auxiliary qubits can be learned from $2^{O(\log(s^2 2^a)^d)} \log(n)$ copies of the Choi state, improving the $n^{O(\log(s^2 2^a)^d)}$ bound by Nadimpalli et al.
Along the way, we give a new proof of the optimal performance of Classical Shadows based on Pauli analysis. We also strengthen the lower bounds against QAC0 to compute the address function. Finally, we propose an approach to improving the PAC learning upper bounds of AC0 circuits, up to an open question in Fourier analysis. Our techniques are based on Fourier and Pauli analysis, and our learning upper bounds are a refinement of the low degree algorithm.
Authors: Jinge Bao and Francisco Escudero Gutiérrez

Title: Lindblad engineering for quantum Gibbs state preparation under the eigenstate thermalization hypothesis

Abstract: Building upon recent progress in Lindblad engineering for quantum Gibbs state preparation algorithms, we propose a simplified protocol that is shown to be efficient under the eigenstate thermalization hypothesis (ETH). The ETH reduces circuit overheads of the Lindblad simulation algorithm and ensures a fast convergence toward the target Gibbs state. Moreover, we show that the realized Lindblad dynamics exhibits an inherent resilience against stochastic noise, opening up the path to a first demonstration on quantum computers. We complement our claims with numerical studies of the algorithm's convergence in various regimes of the mixed-field Ising model. In line with our predictions, we observe a mixing time scaling polynomially with system size when the ETH is satisfied. In addition, we assess the impact of algorithmic and hardware-induced errors on the algorithm's performance by carrying out quantum circuit simulations of our Lindblad simulation protocol with a local depolarizing noise model. This work bridges the gap between recent theoretical advances in dissipative Gibbs state preparation algorithms and their eventual quantum hardware implementation.
Authors: Eric Brunner, Luuk Coopmans, Gabriel Matos, Matthias Rosenkranz, Frederic Sauvage and Yuta Kikuchi

Title: Low Rank Piecewise Polynomial Data Encoding

Abstract: Encoding function values on regular grids is vital for data compression and high-dimensional computations. Exact TT representations exist only for simple functions, and direct polynomial or Fourier expansions yield prohibitive ranks on fine meshes. Hierarchical SVD methods become infeasible as grid resolution grows, and black-box techniques like TT-Cross or multiscale interpolative constructions may fail to converge or incur high computational costs. We propose a novel low-rank TT interpolation scheme to overcome that problem. With this scheme, we can encode noisy data efficiently into a quantum computer with low circuit depth. Numerical experiments in 1D, 2D, and 3D demonstrate high accuracy is achieved with minimal rank growth and lower runtime than re-computing TT-Cross or TT-SVD on the fine grid.
Authors: Siddhartha Emmanuel Morales Guzman

Title: Low-depth measurement-based deterministic quantum state preparation

Abstract: We present a low-depth amplitude encoding method for arbitrary quantum state preparation. Building on the foundation of an existing divide-and-conquer algorithm, we advance this method by proposing how to disentangle the final state. Furthermore, the construction of our

algorithm leads to a natural opportunity to re-use qubits, allowing for flexible, case-specific hybrid algorithms to emerge, before going on to explicitly highlight the potential of our hybrid algorithms in encoding sparse states.
Authors: Roselyn Nmaju, Sarah Croke and Fiona Speirits

Title: Lower Bounding Betti Numbers using Tensor Networks
Abstract: Topological data analysis is applied across a diverse range of fields to extract robust, global features from high-dimensional datasets. However, the computation of these high-dimensional topological features, such as the number of k-dimensional holes (i.e., the k-th Betti number), is challenging, and it is rare in practice to be able to compute topological information beyond k > 2. It was found that the Hodge Laplacian of the homology problem can be mapped to a supersymmetric quantum many-body problem. This connection to many-body physics inspired quantum algorithms to compute Betti numbers. In this work, instead, we present a quantum-inspired tensor network-based method, where Betti numbers are computed as the ground space degeneracy of a fermionic quantum many-body system on a graph. Specifically, we leverage matrix product state methods such as DMRG to probe the dimensionality of the zero-energy eigenstates, resulting in an algorithm using tensor networks to compute lower bounds to the Betti number. We implement our algorithm in Python and perform numerical experiments to demonstrate the correctness of our algorithm up to fourteenth-order Betti numbers. We expect our algorithm to be relevant for higher-order, lower-cardinality Betti numbers.
Authors: Alice Barthe, Casper Gyurik, Jordi Tura, Vedran Dunjko and Patrick Emonts

Title: Many Body Eigenvalue Problems with a Trapped Ion System
Abstract: This work presents a hybrid analog quantum-classical variational algorithm potentially implementable in an ion trap. The cost function, defined from the Hamiltonian, is minimised using classical optimisation techniques to determine optimal variational parameters. This approach is useful for solving eigenvalue problems, particularly in quantum chemistry and materials science. The project demonstrates the integration of quantum evolutions with classical optimisation. The aim is to prepare quantum many-body states at the hardware level as an output rather than just have its classical description. A standard method involves preparing an easily preparable initial state and using adiabatic methods such as quantum annealing, which gradually switch on the Hamiltonian under consideration. However, being slow, this has the issue of decoherence. In this work, non-equilibrium dynamics allows creating the state faster. The Hamiltonian for this non-equilibrium dynamics is created through learning.
Authors: Prashasti Tiwari

Title: Measurement disturbance tradeoffs in unsupervised quantum classification
Abstract: In an increasingly quantum world with more and more quantum technologies nearing practical use, the importance of interacting directly with quantum data is becoming clear. Although doing so often leads to advantages, it also presents us with some uniquely quantum challenges: for example, information about a quantum system cannot, in general, be extracted without disturbing the state of the system. In this presentation, we explore how performing a learning task on quantum data disturbs it, and affects one's ability to learn about it again in the future. In particular, we focus on the learning task of unsupervised binary classification, and how it affects quantum data when it is performed on a subset of it. In such a binary classification task, we are given a dataset that is made up of qubits that are each in one of two unknown pure states, and our aim is

to cluster, with optimal probability of success, the data points into two groups based on their state. To investigate how well we can perform this task sequentially, we first consider a base case of a three-qubit dataset and investigate how an intermediate classification on a two-qubit subset affects the final full classification. We find that there is an analytical tradeoff between the success rates of the two classifications and that, although the intermediate classification does indeed affect the subsequent one in a non-trivial way, there is a remarkably large region where the first classification does not force the second away from its optimal probability of success. Motivated by this feature, we go on to investigate whether an intermediate classification can leave a subsequent one unaffected in the more general setting of an n-qubit dataset. Mathematical and numerical hints lead us to conjecture that nothing about the order of the qubits in an (n-1)-qubit dataset can be learnt without affecting a subsequent classification on the full dataset. We show that an immediate consequence of this is that a non-trivial intermediate classification on smaller subsets of n-m qubits (m > 1) will always negatively impact a subsequent one on all n qubits. After discussing two bounds on how successful an intermediate classification of n-1 qubits can be without affecting the following n-qubit one, we conclude with a discussion of caveats and future lines of research.
Authors: Hector Spencer-Wood, John Jeffers and Sarah Croke

Title: MG-Net: Learn to Customize QAOA with Circuit Depth Awareness

Abstract: Quantum Approximate Optimization Algorithm (QAOA) and its variants exhibit immense potential in tackling combinatorial optimization challenges. However, their practical realization confronts a dilemma: the requisite circuit depth for satisfactory performance is problem-specific and often exceeds the maximum capability of current quantum devices. To address this dilemma, we first analyze the convergence behavior of QAOA, uncovering the origins of this dilemma and elucidating the intricate relationship between the employed mixer Hamiltonian, the specific problem at hand, and the permissible maximum circuit depth. Harnessing this understanding, we introduce the Mixer Generator Network (MG-Net), a unified deep learning framework adept at dynamically formulating optimal mixer Hamiltonians tailored to distinct tasks and circuit depths. Systematic simulations, encompassing Ising models and weighted Max-Cut instances with up to 64 qubits, substantiate our theoretical findings, highlighting MG-Net's superior performance in terms of both approximation ratio and efficiency.
Authors: Yang Qian, Xinbiao Wang, Yuxuan Du, Yong Luo and Dacheng Tao

Title: Modeling Quantum Circuit Parameters with Size-Independent Machine Learning Models

Abstract: The Variational Quantum Eigensolvers aim to approximate ground states of electronic systems by optimizing the parameters of a given quantum circuit. Despite recent improvements, larger systems still require a prohibitively large amount of time and resources on current-day hardware, and the optimization process is a significant computational bottleneck. Machine learning models have been suggested to overcome this bottleneck, but so far most work focuses on single instances or fixed-size instances only, therefore lacking the potential to generalize to larger instances. Here, we develop machine learning methods that are independent of the size of the instance and show that it is sufficient to train the model on small instances and still receive physically reasonable behavior in larger systems.
Authors: Korbinian Stein, Davide Bincoletto and Jakob Kottmann

Title: MPS-based Fourier series loading

Abstract: We develop a matrix product state (MPS) based method to load a function represented by a truncated Fourier series onto a quantum device. Both the circuit depth and the gate count scale only linearly with the number of qubits. Our approach does not require long-range gates, auxiliary qubits, or postselection, in contrast to the recently proposed methods achieving a similar scaling. We test the algorithm on multivariate Gaussian distributions on IBM quantum devices.

Authors: Illia Lukin, Vladyslav Bohun, Mykola Lukhanko and Maciej Koch-Janusz

Title: Neural network-assisted quantum compilation

Abstract: Efficient quantum computation relies heavily on quantum circuit compilation; however, conventional methods produce circuits whose complexity scales exponentially with system size, significantly hindering practical quantum advantage in the NISQ (noisy intermediate-scale quantum) era. We propose a hybrid neural network-assisted quantum circuit compiler combining classical neural networks with variational quantum circuit learning, achieving exponential speedups and substantial reductions in quantum resource usage. We develop a comprehensive mathematical framework for this hybrid approach, validated through successful proof-of-principle and scalable simulations for up to 6-qubit interactions. Extensive benchmarking confirms that our neural network-assisted compiler significantly outperforms state-of-the-art methods, accelerating practical quantum algorithm and simulation development and enabling tangible quantum advantages. This work brings meaningful quantum advantages closer to real-world applicability.

Authors: Wenyu Guo, Qing Liu, Ximing Wang, Chufan Lyu, Jayne Thompson, Andrew J.P. Garner, Aaron Tranter, Rigui Zhou, Chengran Yang and Mile Gu

Title: New perspectives on quantum kernels through the lens of entangled tensor kernels

Abstract: Quantum kernel methods are one of the most popular approaches to quantum machine learning. However, it is not fully understood which kind of problems are suitable for quantum kernels. In this work, we introduce the notion of entangled tensor kernels by generalizing the concept of product kernels, and show that all embedding quantum kernels are entangled tensor kernels. We discuss how this novel perspective allows one to gain insights into both the unique inductive bias of quantum kernels and potential methods for their dequantization.

Authors: Seongwook Shin, Ryan Sweke and Hyunseok Jeonga

Title: Nonlinear Quantum Image Encoding via Information Mixing

Abstract: Efficient quantum image encoding is essential for visual information processing in Quantum Machine Learning (QML) tasks. There are existing methods such as the Flexible Representation of Quantum Images, which use probability amplitude–based encoding of pixel intensities via controlled single-qubit rotations, and Novel Enhanced Quantum Representation (NEQR), which uses direct computational-basis encoding of grayscale pixel values via multi-qubit registers. While these schemes excel at capturing spatial and intensity information, QML models can still struggle to extract and fully exploit all of these encoded features in downstream learning tasks. Inspired by the classical token-mixing approach of FNet, which leverages Fourier transforms to mix token representations, and by Quixer's quantum transformer implementation utilizing Linear Combination of Unitaries (LCU) and Quantum Singular Value Transformation (QSVT), in this work, we introduce an information mixer that interleaves pixel amplitudes via an entangling quantum circuit and injects nonlinearity through QSVT, effectively acting as a quantum "activation" within the encoding process. We integrate this mixer into FRQI and NEQR encodings and evaluate

with quantum neural networks (QNNs) on the MNIST dataset. Compared to unenhanced baselines, these information-mixed QNNs converge faster and achieve better classification scores. Our results underscore the potential of quantum signal processing techniques to advance QML pipelines.
Authors: Natchapol Patamawisut and Yen-Jui Chang

Title: On Bounding Quantum Fidelity with Conformal Prediction

Abstract: Hardware noise makes quantum devices unreliable and limits their applicability to real-world tasks. The theoretical noise characterization required by Quantum Error Correction and Mitigation models is challenging and often impossible to verify.

Quantifying the effects of hardware disturbances via the Bhattacharyya Coefficient between the noisy and noiseless distributions obtained from the same circuit, we show how to bound a device's reliability with conformal prediction, a popular machine learning technique for assumption-free uncertainty quantification. The obtained quantum fidelity lower bounds are independent of the device structure and can be extrapolated to the scenario where test devices are classically simulable (under mild assumptions).

Authors: Nicolo Colombo and Thomas Gargan

Title: On the structure of easy and hard-to-learn positive MPOs

Abstract: In this work we address the problem of quantum noise characterization, a critical task for the design of reliable quantum devices. Noise can be modeled as non-unitary quantum channels, which, for short-depth circuits made of local operations, can be represented efficiently via tensor networks with controlled bond dimension whose parameters can be learned from experimental data.

While efficient algorithms for learning matrix product states exist, the learning of mixed states and quantum channels via their purified representations has proven to be challenging due to the non-convex nature of the optimization landscape, the NP-hardness and undecidability of checking the positivity of a matrix product operator (MPO), and the difficulty in bounding the bond dimension of the purification.

Here we look at the problem following two complementary approaches: first we study the complexity of probably-approximately-correct (PAC)-learning positive semi-definite (psd) MPOs representing quantum states and quantum channels, showing that the two problems are NP-hard. Second, we rely on positivity guarantees given by Petz-reconstructed states in order to learn short-depth tensor representations of channels from the marginals of their Choi matrix. We provide the sample complexity required for this recovery and give upper bounds on the precision of the reconstruction in terms of the infidelity and trace distance of the collected marginals, offering useful practical insights for quantum process tomography.

Authors: Rebecca Erbanni, Gregory A. L. White, Jens Eisert and Matthias C. Caro

Title: Optically Probing Quantum Reservoir Memory

Abstract: Quantum reservoir computing (QRC) presents an innovative framework for leveraging quantum systems in machine learning applications, particularly suited for the noisy intermediate-scale quantum devices era. Traditional metrics such as short-term memory capacity (STMC) are commonly used to assess QRC performance but fall short in elucidating the physical processes underpinning these systems. In this study, we establish a quantitative link between the optical absorption spectrum of a quantum reservoir and its memory capabilities. Our findings demonstrate that optimal STMC coincides with peak absorption, offering a physical rationale for

the previously observed “sweet-spot” behavior in QRC performance relative to dissipation. This relationship integrates quantum information theory with experimentally measurable physical attributes, paving the way for the design of quantum reservoir computers with tailored, task-specific memory capabilities.
Authors: Niclas Götting, Steffen Wilksen, Alexander Steinhoff, Frederik Lohof and Christopher Gies

Title: Optimization Driven Quantum Circuit Reduction

Abstract: Implementing a quantum circuit on specific hardware with a reduced available gate set is often associated with a substantial increase in the length of the equivalent circuit. This process is also known as transpilation, and due to decoherence, it is mandatory to keep quantum circuits as short as possible without affecting functionality. In this work we propose three different transpilation approaches, based on a localized term-replacement scheme, to substantially reduce circuit lengths while preserving the unitary operation implemented by the circuit. The first variant is based on a stochastic search scheme, and the other variants are driven by a database retrieval scheme and a machine learning-based decision support. We show that our proposed methods generate short quantum circuits for restricted gate sets, superior to the typical results obtained by using different Qiskit optimization levels. Our method can be applied to different gate sets and scales well with an arbitrary number of qubits.
Authors: Bodo Rosenhahn, Tobias J. Osborne and Christoph Hirche

Title: Optimizer-Dependent Generalization Bound for Quantum Neural Networks

Abstract: Quantum neural networks (QNNs) play a pivotal role in addressing complex tasks within quantum machine learning, analogous to classical neural networks in deep learning. Ensuring consistent performance across diverse datasets is crucial for understanding and optimizing QNNs in both classical and quantum machine learning tasks, but remains a challenge as QNNs' generalization properties have not been fully explored. In this paper, we investigate the generalization properties of QNNs through the lens of learning algorithm stability, circumventing the need to explore the entire hypothesis space and providing insights into how classical optimizers influence QNN performance. By establishing a connection between QNNs and quantum combs, we examine the general behaviors of QNN models from a quantum information theory perspective. Leveraging the uniform stability of the stochastic gradient descent algorithm, we propose a generalization error bound determined by the number of trainable parameters, data uploading times, dataset dimension, and classical optimizer hyperparameters. Numerical experiments validate this comprehensive understanding of QNNs and align with our theoretical conclusions.
Authors: Chenghong Zhu, Hongshun Yao, Yingjian Liu and Xin Wang

Title: Photonic Quantum Kernel methods for Malware Classification

Abstract: One of the main difficulties of QML methods is the phenomenon of barren plateaus, where the gradients of the loss functions vanish exponentially fast. Some people avoid the barren plateau issue by using quantum kernel methods, which can describe nearly every supervised learning task involving a quantum circuit. For this kind of method, turning to quantum photonic circuits also allows the use of shallow models and control over their expressivity with the input photon number. In this paper, we apply this knowledge about QML and photonic circuits to a concrete problem: malware detection. This is an important cybersecurity topic, and we believe that QML can handle our datasets, which are much more complex than the usual datasets (MNIST,

Iris, etc.), more efficiently than classical models, which would require many samples and features to train correctly.
Authors: Benjamin Stott, Grégoire Barrué and Tony Quertier

Title: Physics-inspired Generative AI models via real hardware-based noisy quantum diffusion

Abstract: Quantum Diffusion Models (QDMs) are an emerging paradigm in Generative AI that aims to use quantum properties to improve the performance of their classical counterparts. However, existing algorithms are not easily scalable due to the limitations of near-term quantum devices. Following our previous work on QDMs, we propose and implement two physics-inspired protocols. In the first, we use the formalism of quantum stochastic walks, showing that a specific interplay of quantum and classical dynamics in the forward process produces statistically more robust models, generating sets of MNIST images with lower Fréchet Inception Distance (FID) than using totally classical dynamics. In the second approach, we realize an algorithm to generate images by exploiting the intrinsic noise of real IBM quantum hardware with only four qubits. Our work could be a starting point to pave the way for new scenarios for large-scale algorithms in quantum Generative AI, where quantum noise is neither mitigated nor corrected, but instead exploited as a useful resource.
Authors: Marco Parigi, Stefano Martina, Francesco Aldo Venturelli and Filippo Caruso

Title: Positive-Unlabeled Learning for Training an Entanglement Detector

Abstract: Entanglement detection, the process of verifying quantum entanglement, is a fundamental challenge in quantum information processing. Various approaches have been proposed to address this challenge, with many recent studies applying supervised machine learning methods. While these methods have demonstrated high accuracy in entanglement detection, it is reasonable to assume that the entangled states themselves are not definitively known. To address this limitation, we have devised a machine learning method for entanglement detection based on positive-unlabeled learning, a classical machine learning framework that does not use label information from negative data. Using a deep neural network model on a synthetic dataset under the assumption of mixed states, we conducted experiments on a classical computer to validate the effectiveness and characteristics of the proposed method. Our approach introduces a novel framework that accounts for the data generation constraints in the training process of an entanglement detector, thereby advancing machine learning techniques in quantum information science.
Authors: Taisei Nohara, Itsuki Noda and Satoshi Oyama

Title: Potential of multi-anomalies detection using quantum machine learning

Abstract: Maintenance of production equipment is critical in manufacturing. Typically, machine learning models are trained on sensor data closely attached to equipment. However, as the number of machines increases, computational cost grows rapidly. In practice, anomalies are often identified by human operators through auditory perception, relying heavily on experience and intuition. In vibration analysis, especially, AR model coefficients combined with one-class SVMs are used for detecting anomalies. In this work, we explore the effect of substituting the classical kernel in the one-class SVM with a quantum kernel. Two experimental setups were used. The first involved a miniature racing car track, where the car passes over a patch of hook-and-loop fastener to generate abnormal sounds, which are recorded using a microphone. The second involved an open-belt drive, where chopsticks are inserted at specific times to produce crushing sounds,

simulating sudden anomalies. Our results show a clear advantage of quantum kernels over classical Gaussian (RBF) kernels. On the miniature car track dataset, the quantum kernel achieved an accuracy and F1-score of 0.82, compared to 0.64 and 0.39 respectively for the RBF kernel. For the crushing device, the quantum kernel achieved perfect accuracy and F1-score (1.00), while the RBF kernel reached only 0.64 accuracy and 0.43 F1-score. These findings suggest that quantum kernels enhance the classification accuracy for diverse types of abnormal sound patterns, including both periodic and impulsive anomalies.
Authors: Takao Tomono and Kazuya Tsujimura

Title: QCA-MolGAN: Quantum Circuit Associative Molecular GAN with Multi-Agent Reinforcement Learning

Abstract: Navigating the vast chemical space of molecular structures to design novel drug molecules with desired target properties remains a central challenge in molecular drug discovery. Recent advances in deep generative models offer promising solutions. This work presents a novel quantum circuit Born machine (QCBM)-enabled Generative Adversarial Network (GAN), called QCA-MolGAN, for generating drug-like molecules. The QCBM serves as a learnable prior distribution, which is associatively trained to define a latent space aligning with high-level features captured by the GAN's discriminator. Additionally, we integrate a novel multi-agent reinforcement learning network to guide molecular generation with desired targeted properties, optimizing key metrics such as quantitative estimate of drug-likeness (QED), octanol-water partition coefficient (LogP), and synthetic accessibility (SA) scores in conjunction with one another. Experimental results demonstrate that our approach enhances the property alignment of generated molecules, with the multi-agent reinforcement learning agents effectively balancing chemical properties.
Authors: Aaron Thomas, Yu-Cheng Chen, Hubert Valencia, Sharu Jose and Ronin Wu

Title: Q-Compression: Quantum-Aware Model Compression Techniques for Scalable Quantum Machine Learning

Abstract: Near-term quantum processors impose strict limits on circuit depth, gate count, and controllable parameters, constraining the deployment of expressive Quantum Machine Learning (QML) models and intensifying barren-plateau effects. We present Q-Compression, a quantum-native framework that systematically reduces circuit complexity while maintaining task accuracy through three complementary operations: (i) entanglement-aware pruning, which removes gates whose marginal contribution to bipartite entanglement entropy and state fidelity is below a rigorously derived threshold; (ii) quantum Fisher information–guided parameter quantization, which snaps rotation angles with low Fisher curvature to coarse discrete values, thereby lowering control precision requirements and mitigating noise sensitivity; and (iii) gradient-variance unitary freezing, which deactivates gate blocks exhibiting vanishing gradient dispersion, shortening effective depth without impairing expressivity. Each transformation is governed by closed-form fidelity bounds and followed by automatic recompilation to the target device's coupling graph and native gate set, ensuring hardware compatibility under realistic noise models. Empirical evaluations focus on multiple benchmark image classification datasets, including MNIST and its established variants, to quantify gate count reduction, depth savings, and accuracy retention on representative tasks. By offering a principled path to resource-efficient QML that aligns with the capabilities of today's noisy intermediate-scale quantum (NISQ) hardware, Q-Compression seeks to advance practical QML applications and provide a unified strategy for resource-efficient QML on NISQ devices and a scalable blueprint for future fault-tolerant architectures.

Authors: Nouhaila Innan and Muhammad Shafique

Title: QKAN: Quantum Kolmogorov-Arnold Networks

Abstract: The potential of learning models in quantum hardware remains an open question. Yet, the field of quantum machine learning persistently explores how these models can take advantage of quantum implementations. Recently, a new neural network architecture, called Kolmogorov-Arnold Networks (KAN), has emerged, inspired by the compositional structure of the Kolmogorov-Arnold representation theorem. In this work, we design a quantum version of KAN called QKAN. Our QKAN exploits powerful quantum linear algebra tools, including quantum singular value transformation, to apply parameterized activation functions on the edges of the network. QKAN is based on block-encodings, making it inherently suitable for direct quantum input. Furthermore, we analyze its asymptotic complexity, building recursively from a single layer to an end-to-end neural architecture. The gate complexity of QKAN scales linearly with the cost of constructing block-encodings for input and weights, suggesting broad applicability in tasks with high-dimensional input. QKAN serves as a trainable quantum machine learning model by combining parameterized quantum circuits with established quantum subroutines. Lastly, we propose a multivariate state preparation strategy based on the construction of the QKAN architecture.
Authors: Petr Ivashkov, Po-Wei Huang, Kelvin Koor, Lirandë Pira and Patrick Rebentrost

Title: Quantitative convergence of trained quantum neural networks to a Gaussian process

Abstract: We study quantum neural networks where the generated function is the expectation value of the sum of single-qubit observables across all qubits. In previous work, it is proven that the probability distributions of such generated functions converge in distribution to a Gaussian process in the limit of infinite width for both untrained networks with randomly initialized parameters and trained networks. In our work, we provide a quantitative proof of this convergence in terms of the Wasserstein distance of order 1. First, we establish an upper bound on the distance between the probability distribution of the function generated by any untrained network with finite width and the Gaussian process with the same covariance. This proof utilizes Stein's method to estimate the Wasserstein distance of order 1. Next, we analyze the training dynamics of the network via gradient flow, proving an upper bound on the distance between the probability distribution of the function generated by the trained network and the corresponding Gaussian process. This proof is based on a quantitative upper bound on the maximum variation of a parameter during training. This bound implies that for sufficiently large widths, training occurs in the lazy regime, i.e., each parameter changes only by a small amount. While the convergence result of previous work holds at a fixed training time, our upper bounds are uniform in time and hold even as time approaches infinity.
Authors: Anderson Melchor Hernandez, Filippo Girardi, Davide Pastorello and Giacomo De Palma

Title: Quantum Bayes' rule from a minimum change principle: gradient-free belief updates for quantum learning

Abstract: In classical machine learning, Bayes' rule underlies principled belief updates based on new evidence. In this work, we extend this logic to quantum systems by introducing a quantum analog of the minimum change principle: a variational approach to inference that minimizes deviation from prior beliefs while remaining consistent with new data. We show that, when applied to quantum channels and states, this principle leads to a closed-form update rule that

coincides with the Petz transpose map, a central structure in quantum information theory. Crucially, this update rule is implementable via quantum algorithms, which may open new possibilities in gradient-free parameter updates in online learning, quantum Bayesian learning, and adaptive quantum protocols.
Authors: Anderson Ge Bai, Francesco Buscemi and Valerio Scarani

Title: Quantum Circuit Optimization for Variational Grover's Search via ZX Calculus

Abstract: Grover’s algorithm is widely regarded as a promising quantum search algorithm that can achieve a quadratic speedup over classical counterparts. However, applying Grover’s search algorithm to many practical problems often requires formulating the problem as a Boolean function and encoding it into an explicit oracle, which is frequently intractable or even impossible for complex domains. To address this limitation, variational Grover's search replaces the fixed oracle with a parameterized (variational) quantum circuit that learns the underlying structure of the search problem from data. By optimizing the variational oracle’s parameters via a classical optimizer, one can approximate the marking function and still leverage Grover-style amplitude amplification to concentrate amplitude on candidate solutions. In this work, we focus on quantum circuit optimization for variational Grover's search using the ZX calculus, a graphical calculus for reasoning about and simplifying quantum circuits. We demonstrate that ZX-based simplification techniques can reduce the overall gate count of the variational oracle and associated diffusion operators, outperforming quantum circuit optimization from Qiskit’s transpiler. Finally, we characterize how gate counts for variational Grover's search circuits, with and without ZX-calculus optimization, scale as problem size and variational circuit depth increase, showing that ZX-optimized implementations grow more slowly than their naively compiled counterparts.
Authors: Natchapol Patamawisut and Ruchipas Bavontaweepanya

Title: Quantum circuits as a game: A reinforcement learning agent for quantum compilation and its application to reconfigurable neutral atom arrays

Abstract: We introduce the quantum circuit daemon (QC-Daemon), a reinforcement learning agent for compiling quantum device operations aimed at efficient quantum hardware execution. We apply QC-Daemon to the move synthesis problem called the Atom Game, which involves orchestrating parallel circuits on reconfigurable neutral atom arrays. In our numerical simulation, the QC-Daemon is implemented by two different types of transformers with a physically motivated architecture and trained by a reinforcement learning algorithm. We observe a reduction of the logarithmic infidelity for various benchmark problems up to 100 qubits by intelligently changing the layout of atoms. Additionally, we demonstrate the transferability of our approach: a Transformer-based QC-Daemon trained on a diverse set of circuits successfully generalizes its learned strategy to previously unseen circuits.
Authors: Kouhei Nakaji, Jonathan Wurtz, Haozhe Huang, Luis Mantilla Calderon, Karthik Panicker, Elica Kyoseva and Alán Aspuru-Guzik

Title: Quantum Deep Learning Force Field

Abstract: Machine learning force fields are emerging as powerful tools for predicting atomic interactions with near density functional theory (DFT) accuracy at significantly reduced computational cost. In this study, we propose a quantum deep learning force field (QDFF) model that leverages quantum neural networks (QNNs) or quantum convolutional neural networks (QCNNs) to represent atomic environments and predict interatomic forces. These models are

constructed from parameterized quantum circuits, where the model parameters are encoded through the rotation angles of quantum gates. We evaluate QDFF on a silicon crystal by comparing the predicted forces and total forces acting on atoms with results from classical machine learning force fields and DFT calculations. Furthermore, we use the QNN or QCNN-predicted forces to compute the phonon dispersion relation and compare them with DFT calculations.
Authors: Hoang Anh Nguyen, Nhu Duc Dinh, Le Tu Uyen Tu and Van Duy Nguyen

Title: Quantum Ensemble Learning with QRAM-Based Subsampling and Shallow Clustering Weak Learners

Abstract: Ensemble learning strategies are widely used in classical machine learning to improve robustness and generalization by reducing prediction variance. Motivated by this, we propose a fully quantum bagging framework, termed Quantum Bootstrapped Bagging (QBB), that exploits quantum subsampling and shallow quantum clustering models. The method relies on quantum random access memory (QRAM) to access training samples in coherent superposition and prepare diverse quantum bootstraps for training. Each quantum bootstrap is provided to a weak learner implemented via a quantum-enhanced k-means clustering algorithm. The final prediction is derived by aggregating results through quantum majority voting (classification) or arithmetic averaging (regression). Experimental evaluations on classification and regression benchmarks validate the performance gains of the proposed method in terms of stability and variance reduction, compared to both classical bagging and single quantum models. Quantum machine learning (QML) offers a promising direction for constructing ensemble-based learning frameworks that operate on quantum-encoded data and exploit diversity among multiple quantum hypothesis states. However, individual quantum models suffer from instability due to stochastic sampling, hardware noise, and training variability. Ensemble learning, particularly bagging, has proven effective in reducing such variance in classical settings. Our goal is to realize a fully quantum analogue of bagging that preserves the computational benefits of QML while achieving the variance-reducing strengths of ensemble methods. We propose Quantum Bootstrapped Bagging, a framework where quantum learners are trained on bootstrapped quantum subsamples accessed via QRAM. This approach ensures diversity among the learners and facilitates fully quantum end-to-end inference. The bagging is applied at the state preparation level, using quantum subsamples drawn from the encoded dataset in superposition, making the entire workflow inherently quantum. The core motivation is that different learners trained on coherent bootstraps explore different hypotheses in Hilbert space, improving ensemble diversity. This variance reduction is quantified empirically across tasks.
Authors: Neeshu Rathi and Sanjeev Kumar

Title: Quantum Feature Maps for High Frequency Time Series

Abstract: In this work, we propose and compare quantum computing architectures for embedding high-frequency time series data into high-dimensional quantum feature spaces. Effective embeddings are essential for many machine learning tasks that are challenging to solve in the original feature space. A relevant feature embedding can enhance inductive inference by improving the separability of examples and capturing complex static and dynamic dependencies, such as similarities, correlations, and causal relationships. In the context of high-frequency financial time series, embeddings need to address critical statistical properties, including long- and short-term memory, non-stationarity, and seasonality. We study two distinct quantum feature map architectures: a recurrent, state evolution-based architecture and a parallel, feedforward architecture.

The recurrent approach leverages the dynamics of open quantum systems or quantum channels to map time series into the space of quantum density operators. The observed time series is generated through discrete-time measurements performed on the emission system. In parallel, we explore a feedforward approach that encodes time series by simultaneously embedding values at different lags into a quantum state requiring no training of quantum circuit parameters.
Through extensive empirical evaluation, we demonstrate the ability of the quantum feature maps to encode intricate temporal patterns and improve the performance of downstream machine learning models. Our results highlight the potential of quantum computing to tackle the unique challenges posed by high-frequency time series data, offering a promising direction for future research in quantum machine learning and financial analytics.
Authors: Vladimir Rastunkov, Vanio Markov, Charlee Stefanski and Daniel Fry

Title: Quantum Large Language Models via Tensor Network Disentanglers

Abstract: We propose a method to enhance the performance of Large Language Models (LLMs) by integrating quantum computing and quantum-inspired techniques. Specifically, our approach involves replacing the weight matrices in the Self-Attention and Multi-layer Perceptron layers with a combination of two variational quantum circuits and a quantum-inspired tensor network, such as a Matrix Product Operator (MPO). This substitution enables the reproduction of classical LLM functionality by decomposing weight matrices through the application of tensor network disentanglers and MPOs, leveraging well-established tensor network techniques. By incorporating more complex and deeper quantum circuits, along with increasing the bond dimensions of the MPOs, our method captures additional correlations within the quantum-enhanced LLM, leading to improved accuracy beyond classical models while maintaining low memory overhead.
Authors: Borja Aizpurua, Saeed Jahromi, Sukhbinder Singh and Román Orús

Title: Quantum machine learning advantages beyond hardness of evaluation

Abstract: Recent years have seen rigorous proofs of quantum advantages in machine learning, particularly when data is labeled by cryptographic or inherently quantum functions. These results typically rely on the infeasibility of classical polynomial-sized circuits to evaluate the true labeling function. While broad in scope, these results however reveal little about advantages stemming from the actual learning process itself. This motivates the study of the so-called identification task, where the goal is to just identify the labeling function behind a dataset, making the learning step the only possible source of advantage. The identification task also has natural applications, which we discuss. Yet, such identification advantages remain poorly understood. So far they have only been proven in cryptographic settings by leveraging random-generatability, the ability to efficiently generate labeled data. However, for quantum functions this property is conjectured not to hold, leaving identification advantages unexplored. In this work, we provide the first proofs of identification learning advantages for quantum functions under complexity-theoretic assumptions. Our main result relies on a new proof strategy, allowing us to show that for a broad class of quantum identification tasks there exists an exponential quantum advantage unless BQP is in a low level of the polynomial hierarchy. Along the way we prove a number of more technical results including the aforementioned conjecture that quantum functions are not random generatable (subject to plausible complexity-theoretic assumptions), which shows a new proof strategy was necessary. These findings suggest that for many quantum-related learning tasks, the entire learning process—not just final evaluation—gains significant advantages from quantum computation.
Authors: Riccardo Molteni, Simon C. Marshall and Vedran Dunjko

Title: Quantum Memory Resource Advantage in Reinforcement Learning

Abstract: Quantum machine learning has attracted growing interest for its potential to outperform classical methods by exploiting quantum effects. In reinforcement learning (RL), most quantum approaches focus on accelerating learning. However, an important question remains: can the use of qubits reduce the working memory requirements in RL tasks? Here, we present RL tasks where agents with qubits have a provable advantage over classical counterparts in efficiently encoding past information that is essential for selecting optimal actions. This working memory resource advantage can be clearly illustrated by the memory gap and reward gap between the use of qubit and cbit in a representative measurement game. We also identify a task that exhibits an infinite-to-constant memory gap. These results demonstrate the existence of RL tasks where the use of quantum memory outperforms classical and suggest that memory resources may be an alternative evaluation of quantum RL agents performance.

Authors: Hon Wai Lau, Aoi Hayashi, William John Munro, Jayne Thompson and Mile Gu

Title: Quantum Neural Networks in Practice: A Comparative Study with Classical Models from Standard Data Sets to Industrial Images

Abstract: Image classification tasks are among the most prominent examples that can be reliably solved by classical machine learning models. In this study, we compare the performance of randomized classical and quantum neural networks (NNs) as well as classical and quantum-classical hybrid convolutional neural networks (CNNs) for the task of binary image classification. We employ two distinct methodologies: (1) using randomized (quantum) NNs on dimensionality-reduced data, and (2) applying (hybrid) CNNs to full image data. We evaluate these approaches on three data sets of increasing complexity: (i) an artificial hypercube dataset, (ii) MNIST handwritten digits (0's and 1's), and (iii) real-world industrial images from laser cutting machines. We analyze correlations between classification accuracy and quantum model hyperparameters, including the number of trainable parameters, feature encoding methods, circuit layers, entangling gate type and structure, gate entangling power, and measurement operators. This provides a systematic hyperparameter optimization study rarely conducted in quantum machine learning literature. For random quantum NNs, we compare their performance with three literature models. Quantitatively, classical and quantum/hybrid models achieved statistically equivalent classification accuracies across most datasets, with no single approach demonstrating consistent superiority. We observe that quantum models show lower variance with respect to initial training parameters, suggesting better training stability. Among the hyperparameters analyzed, only the number of trainable parameters showed a consistent positive correlation with the model performance. Around 94% of the best-performing quantum NNs had entangling gates, although for hybrid CNNs, models without entanglement performed equally well but took longer to converge. Cross-dataset performance analysis revealed limited transferability of quantum models between different classification tasks. The inconsistent importance of entanglement across our experiments highlights the current limited theoretical understanding of how quantum models actually function and what determines their performance. Our study provides an industry perspective on quantum machine learning for practical image classification tasks, highlighting both current limitations and potential avenues for further research in quantum circuit design, entanglement utilization, and model transferability across varied applications. Text and talk based on the preprint arXiv:2411.19276.

Authors: João F. Bravo, Daniel Basilewitsch, Christian Tutschku and Frederick Struckmeier

Title: Quantum neuromorphic computing with parametrically coupled bosonic modes in superconducting resonators

Abstract: Bosonic modes offer several advantages for quantum neural networks: when neurons are encoded in Fock states, the accessible Hilbert space grows faster than in qubit-based systems; they enable simultaneous coupling via multiple parametric processes; and they are expected to be less susceptible to barren plateaus during training. We explore these features both experimentally using superconducting circuits and through simulations, aiming to assess whether bosonic systems indeed mitigate barren plateaus, whether quantum coherence contributes to learning, and what data encoding strategies are most efficient in quantum systems. Experimentally, we implement quantum reservoir computing with a single bosonic mode in a superconducting resonator, coupled to a transmon qubit for readout. We demonstrate that classical classification tasks can be solved with significantly fewer measured features compared to classical neural networks. We also find that the combination of projective measurement nonlinearity and intrinsic Kerr nonlinearity improves performance. In simulations, we show that training the parametric couplings between bosonic modes via backpropagation increases the model's expressivity. Using a network of six coupled bosonic modes, we classify the full DIGITS dataset, and demonstrate that training further reduces the number of measurements required, compared to the untrained reservoir approach.

Authors: Danijela Markovic

Title: Quantum reservoir computing for temporal series processing using spin-boson systems

Abstract: Quantum reservoir computing is a quantum analog of classical reservoir computing, and is mainly used for classification or prediction of time-series. The reservoir is a complex quantum system that is dynamically driven by the inputs. A set of observables is measured in the output layer and mapped to the target output using linear regression, while the reservoir's degrees of freedom remain unchanged. The main advantage compared to variational quantum algorithms is that the training of QRC does not suffer from barren plateau or local minima. Spin-boson systems are an interesting choice for the reservoir. They couple spin systems to infinite-level harmonic oscillators, and can be easily simulated experimentally. Previous works use these systems for time-series prediction, however they lack a physical insight into the origin of nonlinearity in these systems, as well as a systematic benchmarking of the reservoir's memory, which is an important aspect of QRC. In our work, we use the Jaynes-Cummings (JC) model as the reservoir, and study the emergence of nonlinearity in these systems. We quantify the reservoir' memory using standard linear and nonlinear memory tasks. We study the dynamical features of the observables with respect to varying system parameters, which empirically gives an idea of the parameter regime necessary for the reservoir to perform well. Furthermore, we examine the performance of JC model, as well as JC model in dispersive limit, for complex chaotic time-series forecasting. Our work provides a clear illustration of QRC using these systems.

Authors: Sreetama Das, Gian Luca Giorgi and Roberta Zambrini

Title: Quantum reservoir computing with multiple photonic memristors and skip connections

Abstract: Reservoir computing has recently got significant attention in the context of quantum machine learning, with progress in both numerical results and experimental implementations. In our recent work (Selimović et al., 2025), we demonstrated the first experimental use of single-photon quantum memristors for nonlinear task prediction. Here, we present a novel architecture that combines several quantum memristors in a series-connected setup to enhance predictive performance. Furthermore, we introduce the concept of skip connections – borrowed from classical machine learning – to further boost results. The proposed architecture remains

compatible with current experimental capabilities and yields promising numerical outcomes, outperforming the single-memristor configuration.
Authors: Michał Siemaszko, Martin Mauser, Iris Agresti, Philip Walther and Magdalena Stobińska

Title: Quantum Reservoir Computing with Small Datasets

Abstract: Classical machine learning models can struggle with performance when training on small datasets. In this scenario, there is a risk of overfitting and inability to generalise to new data or being driven by outliers. The performance of quantum machine learning (QML) methods on small datasets is therefore a subject of interest. Early results indicate that quantum reservoir computing (QRC) may be more robust than standard classical models as the size of the dataset decreases [1]. We investigate this by performing a binary classification on a dataset of size 108 with 56 features. We make use of QuEra's Aquila machine [2] to implement the quantum reservoir. We use the SHapley Additive exPlanations (SHAP) method [3] in a classical pre-processing step to identify useful features, reducing the number of data features to be passed to the quantum reservoir. We start with a hyperparameter search with six classical machine learning models, evaluating the best set of hyperparameters for each model using accuracy and SHAP for feature importance. We aggregate features based on the combined results. This aggregate feature list helps form the input to the quantum reservoir. We run QRC for 1 to 15 input features, using the method outlined in [4]. Each input feature set produces a quantum output feature set, which are from expectation values of local observables, in particular <Z_i> and <Z_i Z_j>. These expectation values are evaluated at 8 different time-steps, which results in a larger feature set. As a final step, we train our original classical models on the quantum features. The top row of figure 1 outlines the performance of each model. These results indicate that when training an all quantum output features, there is no performance advantage over training with the original features. However, by performing further SHAP analysis on the quantum output features and training on a reduced set of quantum features, it is possible to achieve similar and often greater performance than training on the original features. The bottom row of figure 1 outlines these results.
References: [1] D. Beaulieu et al. 2024. https://doi.org/10.48550/arXiv.2412.06758 [2] J. Wurtz et al. 2023. https://doi.org/10.48550/arXiv.2306.11727 [3] S. Lundberg et al. 2017. https://doi.org/10.48550/arXiv.1705.07874 [4] M. Kornjaca et al. 2024. https://doi.org/10.48550/arXiv.2407.02553
Authors: Luke Antoncich, Milan Kornjaca, Jonathan Wurtz, Jing Chen, Pascal Elahi and Casey Myers

Title: Quantum Simulations of Chemical Reactions: Achieving Accuracy with NISQ Devices

Abstract: Recently, a lot of interest has been generated in the field of computational chemistry to simulate chemical reactions in the realms of quantum computation using Noisy Intermediate Scale Quantum (NISQ) Computers [1]. The hybrid quantum-classical algorithm, Variational Quantum Eigensolver (VQE), has proven to be an efficient tool for simulating molecules and their properties [2]. In VQE, the variational principle, combined with a classical gradient-based optimizer, is used to determine the ground state energy of a molecule [3,4]. However, due to the limitations in qubit count, circuit depth, and noise in the VQE model, simulating molecular properties does not achieve the desired chemical accuracy when compared to traditional computational chemistry methods, even after considering different properties of a reaction (e.g. activation energy) [5]. In this work, we focus on achieving chemical accuracy in calculating reaction energies for various closed-shell reactions using both VQE and computational chemistry techniques. To address this, we leverage the point-group symmetry of each reactant and product in a reaction to identify excitations that share the same symmetry as the ground state energies. As part of our study, we simulated five

different reactions, and calculated their corresponding reaction energies using both computational chemistry and VQE methods. We were able to achieve chemical accuracy for all five reactions up to an order of 10^{-3}, with one exception where the chemical accuracy remains at the order of 10^{-1}. Further, the total number of combinations of qubits and electrons has been reduced to only one instead of 15, 15, 300, 256, 48 for the five reactions respectively.
Authors: Maitreyee Sarkar, Lisa Roy, Aksah Gutal, Atul Kumar and Manikandan Paranjothy

Title: Quantum State Preparation using Dynamical Invariants and Machine Learning

Abstract: Quantum state preparation is an important first step in almost all applications of quantum technology including quantum algorithms, quantum communications, quantum sensing and quantum error correction. However, engineered quantum systems are susceptible to noise, posing a challenge in characterization and control. Here we address the problem of state preparation in the presence of non-Markovian noise which is commonly observed in many physical devices. We present a control algorithm based on a hybrid between physics-based methods, particularly dynamical invariants pulse design, and machine-learning methods. Additionally, we address some of the limitations of the invariant method that might prevent it from being applied in practice, such as unbounded control signal. We show numerical results for preparing an arbitrary state for a qubit subject to multi-axis random telegraph noise. Our results show high fidelity operations and highlight the capability of hybrid physics and machine learning methods.
Authors: Ritik Sareen, Akram Youssry, Yule Mayevsky and Alberto Peruzzo

Title: Quantum variational parameters as classical data outputs

Abstract: Many machine learning frameworks aim to transform samples from one distribution into another, often by minimizing transport cost functions or learning continuous flows that interpolate between source and target data points. In this work, we propose a quantum model that can be viewed as an analogue of flow matching for modeling transport maps, with the key idea to encode and retrieve data as parameters of a quantum circuit. The model uses a parameterized quantum circuit to learn from a dataset of source-target data points related by a map. The circuit takes as input a pair of data points and variational parameters. The latter are optimized with the aim of creating a good energy landscape that connects source to target data points. During the generation stage, the variational parameters are fixed and the data points are initialized to be identical. The optimization is then performed over one of the data points to recover a data pair related by the learned map. This notably allows the generation of classical data without the state tomography commonly used by other approaches. Based on known results, this model can be shown to be able to represent any continuous map. The flexibility of the model is tested in the following tasks: (i) learning a non-linear two-dimensional map, (ii) learning a kernel for image up-scaling, and (iii) learning maps in the latent space of a variational autoencoder. The model was successful in all three tasks while requiring only a modest number of qubits (2, 4, and 4, respectively) and displayed an absence of barren plateaus in the relevant neighborhoods during generation. Notably, as shown in test (iii), this method allows for successfully generating data with dimensions larger than that of the used quantum state at the expense of optimization overhead.Authors: Luis Ernesto Campos Espinoza and Dmitry Guskov

Title: Quantum-Hybrid Siamese Networks with Inter-Channel Weight Interaction

Abstract: Siamese networks are a type of neural network architecture designed to compare two inputs by processing them through identical subnetworks with shared weights and commonly used

for tasks like face verification, signature recognition with the goal is to determine (dis)similarity. However, they are computationally expensive and difficult to train at scale. Hence, this paper proposes a hybrid quantum Siamese network with intension of reducing the number of parameters with the introduction of quantum component, e.g., variational quantum circuit (VQC), inter-channel weight interaction using controlled rotation gates, amplitude embedding. We present a novel quantum-hybrid Siamese architecture that reduces model complexity while maintaining strong performance on signature verification tasks. We explore difference variations of the quantum circuits and combinations between the quantum circuits and the classical models to experiment with the accuracies for different datasets. Inspired from the classical model where weight sharing mechanism is used between the twin networks, we use variational quantum circuit (VQC) that captures similarity between feature embeddings for both the input images (we call these twin networks channels A and B for the two input images). In the classical Siamese models, the model is trained on the basis of the distance between the final outputs of the twin networks. The most commonly used loss function used to train a Siamese model is the contrastive loss function, where the Euclidean distance is used. In our work, we use the overlap between the final states after all the VQC and the inter-channel weight interactions, using a Hadamard-test-based overlap measurement. The model uses amplitude embedding of normalized features, with controlled and non-controlled VQC blocks modulated via a switch qubit to enable channel interaction. Our proposed hybrid models show comparable accuracy than lighter classical baselines while using fewer parameters than the original 6M SigNet model. We also demonstrate a transfer learning setup by freezing classical layers of the 6M SigNet model and optimizing only the quantum circuit. Evaluation of the models on the BHSig260 and CEDAR datasets confirms the expressiveness and parameter efficiency of our design, making it a promising step toward scalable quantum-enhanced machine learning. Our work focuses on achieving accuracies close to the original classical model, but with less model complexity. The experiments show promising results with the best accuracy being 83% when a 2M half-SigNet classical model is combined with one of the proposed quantum circuits in the transfer-learning setup with frozen classical layers.
Authors: Soham Pawar and Dibakar Das

Title: Quantum-Inspired Optimization for High Energy Physics

Abstract: Various computationally challenging tasks in high energy physics can be formulated as quadratic unconstrained binary optimization (QUBO) or Ising problems. This class of problems is designed so that the ground state of the QUBO/Ising Hamiltonian provides the correct answer. Simulated bifurcation (SB) is a promising quantum-inspired approach to solve such problems. SB predicts the ground state by classically emulating the quantum adiabatic evolution of Kerr-nonlinear parametric oscillators, exhibiting bifurcation phenomena to represent the two Ising spin states. Being a quantum-inspired but classical algorithm, it neither suffers from quantum hardware noise nor the data-size limitations that it can handle up to around million-level data size. In contrast to simulated annealing, SB can run in parallel and also benefits from cutting-edge computing resources such as GPUs and FPGAs. I will present its recent applications to charged particle pattern recognition and jet clustering. For track reconstruction, we have observed as much as four orders of magnitude speedup from simulated annealing. SB is also capable of pursuing multiple jet clustering, which is formulated as fully connected QUBO problems that are notoriously known for their difficulty. SB successfully reconstruct multiple jets in one go with the QUBO formulation.
Authors: Hideki Okawa, Qing-Guo Zeng, Xian-Zhe Tao and Man-Hong Yung

Title: Qubit Trajectory Analysis in Quantum Neural Networks

Abstract: In this paper, we analyze the Fourier structure hidden in single-qubit, data-reuploading parameterized quantum circuits (PQCs). We show that each data-encoding rotation forces the qubit to trace a circle whose radius is fixed by the circuit's trainable parameters. Formally, we prove the projected circle theorem and show that the projection of this trajectory onto the measurement axis yields the circuit's Fourier coefficients for a single-layer PQC. The result gives concrete design guidance: select rotation axes that are non-collinear with the encoding axis to avoid degenerate projections. The same geometric insight scales naturally to L-layer circuits, where flexible rotation axes at each layer sustain full expressivity even for targets with rich, high-frequency spectra.

Authors: Seungcheol Oh, Chaemoon Im and Joongheon Kim

Title: Qudit shadow estimation based on the Clifford group and the power of a single magic gate

Abstract: We clarify the sample complexity of qudit shadow estimation based on the Clifford group, where the local dimension d is an odd prime. We show that the overhead of qudit shadow estimation over the qubit counterpart is only O(d), independent of the qudit number n, although the set of stabilizer states may deviate exponentially from a 3-design with respect to the third moment operator. Furthermore, by adding one layer of magic gates, we propose a simple circuit that can significantly boost the efficiency. Actually, a single magic gate can eliminate the O(d) overhead and bridge the gap from the qubit setting.

Authors: Chengsi Mao, Changhao Yi and Huangjun Zhu

Title: Quixer: A Quantum Transformer Model

Abstract: Progress in the realisation of reliable large-scale quantum computers has motivated research into the design of quantum machine learning models. At the same time, large language models, typically implemented using the transformer architecture, have become a cornerstone of natural language processing. Here, we present Quixer: a novel quantum transformer model which utilises the Linear Combination of Unitaries and Quantum Singular Value Transform primitives as building blocks. Quixer operates by preparing a superposition of tokens and applying a trainable non-linear transformation to this mix. We present the first results for a quantum transformer model applied to a practical language modelling task, obtaining results competitive with an equivalent classical baseline.

Authors: Nikhil Khatri, Gabriel Matos, Luuk Coopmans and Stephen Clark

Title: Qutrit-Based Quantum Circuit Design via Reinforcement Learning for Simulating Three-Flavor Collective Neutrino Oscillations

Abstract: Designing quantum circuits capable of simulating complex many-body systems is one of the core challenges in quantum computing. Traditional circuit design methods often rely on heuristic algorithms, which become increasingly inefficient as the system size and complexity grow. In this study, we propose a novel reinforcement learning (RL)-based framework to automate the construction of quantum circuits composed of qutrit gates, focusing on the simulation of collective oscillations in three-flavor neutrino systems. Our approach employs a policy-gradient RL algorithm, where an agent learns to construct sequences of qutrit gates that approximate the unitary time-evolution operator generated by the many-body neutrino Hamiltonian. By using qutrits instead of qubits, the algorithm can more naturally represent the SU(3) algebra governing the system's dynamics, thereby reducing circuit depth and improving expressibility. We validate this

framework through numerical simulations by comparing the learned circuits with theoretical evolution operators. The results show that the RL-designed circuits achieve high fidelity with fewer layers, demonstrating the efficiency and accuracy of the proposed approach.
Authors: Hoang-Anh Nguyen, Duy-Tung Nguyen, Tran-Hien Vo, Dang-Khanh Nguyen, Nhu-Duc Dinh and Van-Duy Nguyen

Title: Ravines in quantum cost landscapes: Opportunities for enhanced VQA predictions on quantum data

Abstract: The geometric and topological structure of the quantum cost landscape (QCL) critically governs the optimization and thus the predictive power of variational quantum algorithms (VQAs). We systematically analyze ravines—low-loss paths connecting local minima—using an adapted version of the nudged elastic band (NEB) algorithm, a method originating from theoretical chemistry. By training VQAs to classify the concentratable entanglement of quantum states, we apply the NEB algorithm to reveal numerical evidence of ravine structures in the QCL with quantum input data. Beyond visualizing these features, we exploit configurations along ravines to construct an ensemble prediction framework. To make ensemble predictions, we average the predictions of multiple models, each using a parameter setting corresponding to a point along the low-loss NEB path. This approach demonstrates compelling performance, combining high classifier independence with robust predictive accuracy of each individual classifier. Notably, our method surpasses classical techniques like Random Forest in test-set performance for suitable choices of hyperparameters. These results provide evidence for ravines as a structural resource in the QCL, offering a strategy to enhance VQA prediction capabilities and thus leveraging quantum landscape topology to improve the reliability and efficacy of VQAs.
Authors: Felix J. Beckmann and João F. Bravo

Title: Real classical shadows

Abstract: Efficiently learning expectation values of a quantum state using classical shadow tomography has become a fundamental task in quantum information theory. In a classical shadows protocol, one measures a state in a chosen basis W after it has evolved under a unitary transformation randomly sampled from a chosen distribution U. In this work we study the case where U corresponds to either local or global orthogonal Clifford gates, and W consists of real-valued vectors. Our results show that for various situations of interest, this “real” classical shadow protocol improves the sample complexity over the standard scheme based on general Clifford unitaries. For example, when one is interested in estimating the expectation values of arbitrary real-valued observables, global orthogonal Cliffords decrease the required number of samples by a factor of two. More dramatically, for k-local observables composed only of real-valued Pauli operators, sampling local orthogonal Cliffords leads to a reduction by an exponential-in-k factor in the sample complexity over local unitary Cliffords. Finally, we show that by measuring in a basis containing complex-valued vectors, orthogonal shadows can, in the limit of large system size, exactly reproduce the original unitary shadows protocol.
Authors: Maxwell West, Antonio Anna Mele, Martín Larocca and Marco Cerezo

Title: Recurrent Measurement-Based Quantum Machine Learning (MBQML) with Feedback

Abstract: We introduce a recurrent hybrid quantum-classical machine learning model leveraging Measurement-Based Quantum Computing (MBQC). Our model processes sequential data using adaptive quantum circuits informed by classical measurement feedback. The framework integrates

parameterized single-qubit rotations, entangling operations, and dynamically updated biases, capturing temporal correlations over input sequences.
Authors: Abdullah Kazi and Jayesh Hire

Title: Reinforcement Learning Based Quantum Circuit Optimization via ZX-Calculus

Abstract: 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 a few tens of gates, the agent consistently improves the state of the art for this type of circuits, for at least up to 80 qubits and 2100 gates, while remaining competitive in terms of computational performance. Additionally, we illustrate the versatility of the agent by incorporating additional optimization routines into the workflow during training, improving the two-qubit gate count state of the art on multiple structured quantum circuits for relevant applications of much larger dimension and different gate distributions than the circuits the agent trains on. This conveys the potential of tailoring the reward function to the specific characteristics of each application and hardware backend. Our approach is a valuable tool for the implementation of quantum algorithms in the near-term intermediate-scale range (NISQ).
Authors: Jan Nogué Gómez

Title: Role of scrambling, noise, and symmetry in temporal learning with quantum systems

Abstract: Scrambling quantum systems have been demonstrated as effective substrates for temporal information processing. While their potential and limitations for learning static data has been studied in our previous work, a theoretical understanding of their performance in temporal tasks, where quantum systems work as both computational substrate and memory, is still lacking. Going beyond our previous setting of a single step iteration and static data, here we consider a general quantum reservoir processing framework that captures a broad range of temporal learning models using quantum systems. We examine the scalability and memory retention of the model with scrambling reservoirs modelled by high-order unitary designs in both noiseless and noisy settings. In the former regime, we show that measurement readouts become exponentially concentrated with increasing reservoir size, yet strikingly do not worsen with the reservoir iterations. Thus, while repeatedly reusing a small scrambling reservoir with quantum data might be viable, scaling up the problem size deteriorates generalization unless one can afford an exponential shot overhead. In contrast, the memory of early inputs and initial states decays exponentially in both reservoir size and reservoir iterations. In the noisy regime, we also prove exponential memory decays with iterations for local noisy channels. We further prove that the scalability limitations also hold for model variants that incorporate mid-circuit measurements and feed-forward, which have been demonstrated as a source that improves performance of learning models. To sidestep these scalability barriers, we show that by employing less scrambling or symmetrized reservoirs, the exponential concentration of outputs can be avoided. Proving these results required us to introduce new proof techniques for bounding concentration in temporal quantum learning models. Our results lay the ground work for understanding the potential of a quantum advantage for temporal learning tasks using quantum reservoir processing.
Authors: Weijie Xiong, Zoe Holmes, Armando Angrisani, Yudai Suzuki, Thiparat Chotibut and Supanut Thanasilp

Title: Sample-Efficient Estimation of Nonlinear Quantum State Functions

Abstract: Efficient estimation of nonlinear functions of quantum states is crucial for various key tasks in quantum computing, such as entanglement spectroscopy, fidelity estimation, and feature analysis of quantum data. Conventional methods using state tomography and estimating numerous terms of the series expansion are computationally expensive, while alternative approaches based on a purified query oracle impose practical constraints. In this article, we introduce the quantum state function (QSF) framework by extending the SWAP test via linear combination of unitaries and parameterized quantum circuits. Our framework enables the implementation of arbitrarily normalized degree-n polynomial functions of quantum states with precision ε using O(n/ε^2) copies. We further apply QSF for developing quantum algorithms for fundamental tasks, including entropy, fidelity, and eigenvalue estimations. Specifically, for estimating von Neumann entropy, quantum relative entropy, and quantum state fidelity, where κ and γ represent the minimal nonzero eigenvalue and normalized factor, respectively, we achieve a sample complexity of Õ(γ^2/(ε^2κ)). Our work establishes a concise and unified paradigm for estimating and realizing nonlinear functions of quantum states, paving the way for the practical processing and analysis of quantum data.

Authors: Hongshun Yao, Yingjian Liu, Tengxiang Lin and Xin Wang

Title: Scalable Quantum Architecture Search based on Relative Fluctuation of Landscapes

Abstract: Balancing trainability and expressibility is a central challenge in variational quantum computing, and quantum architecture search (QAS) plays a pivotal role by automatically designing problem-specific parameterized circuits that address this trade-off. In this work, we introduce a scalable, training-free QAS framework that efficiently explores and evaluates quantum circuits through relative fluctuation. This landscape fluctuation captures key characteristics of the cost function landscape, enabling accurate prediction of circuit learnability without costly training. By combining this metric with a streamlined two-level search strategy, our approach identifies high-performance, large-scale circuits with higher accuracy and fewer gates. We further demonstrate the practicality and scalability of our method, achieving significantly lower classical resource consumption compared to prior work.

Authors: Chenghong Zhu, Xian Wu, Hao-Kai Zhang, Sixuan Wu, Guangxi Li and Xin Wang

Title: Self-Optimizing Quantum Circuits via Reinforcement Learning and Language Models

Abstract: This paper introduces a novel approach to quantum circuit optimization (QCO) using reinforcement learning (RL) applied to a language model (LM) pretrained on quantum circuits. Instead of relying on predefined optimization rules or the unification of diverse techniques, we leverage the LM's existing knowledge of quantum circuits to discover novel optimization pathways. The LM acts as an agent trained via RL, with a reward function focused on minimizing the number of gates possessing multiple inputs; this penalty function is the only user-specified parameter, but it can be an arbitrary numerical criterion. The LM learns to identify and apply advantageous gate combinations beyond those implicitly encoded in its pre-training, effectively refining its existing knowledge and discovering novel optimization strategies. This method avoids the need for explicit integration of diverse optimization techniques, offering a more adaptable and extensible solution for future advancements in quantum compilation. We evaluate the performance of our approach on various circuit families, measuring gate count reduction, circuit depth, and fidelity preservation,

and analyze the emergent optimization patterns, demonstrating the ability of the RL process to improve upon the LM's initial capabilities.
Authors: Antonin Sulc

Title: Simulation cost of classically simulable quantum machine learning models
Abstract: Variational Quantum Algorithms (VQAs) and Quantum Machine Learning (QML) models have gained attention for their potential to exploit near-term quantum devices. However, their scalability is highly challenged by the barren plateau phenomenon, wherein the cost function gradients vanish exponentially with system size. While recent advances have introduced barren plateau–free architectures and mitigation strategies, many of these also render the models classically simulable, raising doubts about their ability to achieve quantum advantage. Indeed, a growing number of classical surrogates have been proposed that aim to replicate the behavior of quantum models on classical devices. Despite this wave in classical approaches, it remains an open and pressing question: Is it computationally more efficient to simulate quantum algorithms classically than to run them on quantum hardware? In this work, we examine this question by analyzing specific barren plateau–free circuit families, each paired with architecture-specific classical surrogates and comparing the computational complexities of Quantum and Classical simulations using three key metrics: classical time, quantum time, and quantum sample complexity. Our analysis provides insight into when classical surrogates remain efficient and when quantum processing may offer a practical advantage, contributing to a deeper understanding of the resource trade-offs in VQAs and QML models.
Authors: Su Yeon Chang, Supanut Thanasilp, Zoë Holmes and Marco Cerezo

Title: Size-Invariant Properties at Depth 1 of the Equivariant Quantum Circuit
Abstract: The Equivariant Quantum Circuit (EQC) has been shown to achieve near-optimal performance in solving small TSP problems (with 20 nodes and fewer) using only two parameters at Depth 1. Since its introduction in 2023, that work has received many citations. Adding to this literature, we offer a new interpretation of the parameters in the EQC. More precisely, we show that effective EQC policies at depth 1 reside within a small parameter space that is independent of the size of the TSP instance. In other words, similar parameter settings across different sizes of the TSP problem result in near-optimal tours that the agent can model at depth 1. We outline our theoretical results and verify them empirically across multiple TSP sizes.
Authors: Jonathan Teo, Xin Wei Lee and Hoong Chuin Lau

Title: Spectral Bias in Parameterised Quantum Circuits
Abstract: In this work, we investigate the phenomenon of spectral bias in quantum machine learning, where, in classical settings, models tend to fit low-frequency components of a target function earlier during training than high-frequency ones, demonstrating a frequency-dependent rate of convergence. We study this effect specifically in parameterised quantum circuits (PQCs). Leveraging the established formulation of PQCs as Fourier series, we prove that spectral bias in this setting arises from the redundancy of the Fourier coefficients, which denotes the number of terms in the analytical form of the model contributing to the same frequency component. The choice of data encoding scheme dictates the degree of redundancy for a Fourier coefficient. We find that the magnitude of the Fourier coefficients' gradients during training strongly correlates with the redundancy of the coefficients. We then further demonstrate this empirically with two different encoding schemes. Additionally, we demonstrate that PQCs with greater redundancy exhibit

increased robustness to random perturbations in their parameters at the corresponding frequencies.
Authors: Callum Duffy, Marcin Jastrzebski and Sarah Malik

Title: Strengthening the no-go theorem for QRNGs
Abstract: Quantum random number generators (QRNGs) are essential for security against quantum algorithms. Randomness as a beacon is a service provided to companies and governments to upgrade their security standards from RSA to PQC-QKD or PQC-RSA protocols. How does an entity ensure that the beacon service has a quantum signature besides relying on faith? Researchers claim that this is indecipherable and have stated a no-go theorem for post-processed bit-streams (Physical Review A 109, 022243 (2024)). In this work, we corroborate the results of the no-go theorem while discussing its nuances using two different random number generators and four test methods.
Authors: Vardaan Mongia, Abhishek Kumar, Shashi Prabhakar and R.P. Singh

Title: Subspace Preserving Quantum Convolutional Neural Network Architectures
Abstract: Subspace preserving quantum circuits are a class of quantum algorithms that, relying on some symmetries in the computation, can offer theoretical guarantees for their training. Those algorithms have gained extensive interest as they can offer polynomial speed-up and can be used to mimic classical machine learning algorithms. In this work, we propose a novel convolutional neural network architecture model based on Hamming weight preserving quantum circuits. In particular, we introduce convolutional layers, and measurement based pooling layers that preserve the symmetries of the quantum states while realizing non-linearity using gates that are not subspace preserving. Our proposal offers significant polynomial running time advantages over classical deep-learning architecture. We provide an open source simulation library for Hamming weight preserving quantum circuits that can simulate our techniques more efficiently with GPU-oriented libraries. Using this code, we provide examples of architectures that highlight great performances on complex image classification tasks with a limited number of qubits, and with fewer parameters than classical deep-learning architectures.
Authors: Léo Monbroussou, Jonas Landman, Letao Wang, Alex Bredariol Grilo and Elham Kashefi

Title: Tensor Network-Enhanced Variational Quantum Circuits: Theory, Hybrid Architectures, and Scalable Optimization for Quantum Machine Learning
Abstract: Variational Quantum Circuits (VQCs) have become central to quantum machine learning on devices with Noisy Intermediate-Scale Quantum (NISQ). Yet, they remain constrained by fundamental issues such as limited representation power, training inefficiencies, and sensitivity to noise. Our work presents a cohesive research agenda based on three synergistic contributions that combine tensor network theory, quantum-classical hybridization, and novel training frameworks. Each approach addresses a different bottleneck in VQC design and training: (1) TTN-VQC: End-to-end quantum learning with theoretical guarantees; (2) VQC-MLPNet: Hybrid quantum-classical architecture with improved expressivity and generalization; (3) Tensor-Guided VQC Training: Classical tensor networks as differentiable controllers of VQC gate parameters. This unified framework enables scalable, trainable, and noise-resilient quantum learning, facilitating practical deployment in science and engineering domains.
Authors: Jun Qi, Chao-Han Huck Yang and Pin-Yu Chen

Title: The Born Ultimatum: Simulability in Quantum Generative Models. Capturing Higher-Order Moments from Data

Abstract: Quantum generative models, like their classical counterparts, are data-driven, and their performance depends on the structure of the underlying distribution. In this work, we examine the Quantum Circuit Born Machine (QCBM) through the framework of k order correlations. Describing the probability distribution in terms of these correlations allows for truncation and surrogate training using tensor networks or Pauli propagation, independently of the chosen loss function. This perspective enables us to move beyond the two-qubit layer ansatz used in past literature that has been restricted to the Maximum Mean Discrepancy loss. We highlight that higher-order correlators can be both significant and trainable. Finally, we take an initial step towards analysing the potential advantage in scaling when the QCBM is trained classically and deployed in a quantum computer, once again by means of k-order correlations.

Authors: Mario Herrero-González, Ross Grassie, Kieran McDowall, Sjoerd Beentjes, Ava Khamseh and Elham Kashefi

Title: The effect of classical optimizers and Ansatz depth on QAOA performance in noisy devices

Abstract: The Quantum Approximate Optimization Algorithm (QAOA) is a variational quantum algorithm for Near-term Intermediate-Scale Quantum computers (NISQ) providing approximate solutions for combinatorial optimization problems. The QAOA utilizes a quantum-classical loop, consisting of a quantum ansatz and a classical optimizer, to minimize some cost functions computed on the quantum device. This presentation presents an investigation into the impact of realistic noise on the classical optimizer and the determination of optimal circuit depth for the Quantum Approximate Optimization Algorithm (QAOA) in the presence of noise. While there is no significant difference in the performance of classical optimizers in a state vector simulation, the Adam and AMSGrad optimizers perform best in the presence of shot noise. Under the conditions of real noise, the SPSA optimizer, Adam, and AMSGrad emerge as the top performers. The study also reveals that the quality of solutions to some five-qubit minimum vertex cover problems increases for up to around six layers in the QAOA circuit, after which it begins to decline. This analysis shows that increasing the number of layers in the QAOA to increase accuracy may not work well in a noisy device.

Authors: Ilya Sinayskiy

Title: The Ubiquity of QPINNs: A Multi-Domain Study

Abstract: Physics-informed neural networks (PINNs) have emerged as a powerful paradigm for solving partial differential equations (PDEs) by embedding physical laws directly into neural network architectures. However, classical PINNs face fundamental limitations when confronting high-dimensional, nonlinear systems with complex boundary conditions and multi-scale dynamics. This work presents a systematic evaluation of Quantum Physics-Informed Neural Networks (QPINNs), demonstrating their universal applicability and computational advantages across three disparate scientific domains that collectively span the breadth of modern computational physics and quantitative science: (1) the Navier-Stokes equations for nonlinear fluid turbulence, (2) Einstein’s field equations describing gravitational dynamics in general relativity, and (3) the Black-Scholes partial differential equation fundamental to financial derivative pricing. Our quantum-enhanced framework leverages parameterized quantum circuits to encode solution spaces in high-dimensional Hilbert spaces, exploiting quantum superposition to represent exponentially complex function approximations that are intractable for classical neural networks.

Across all three domains, we establish rigorous theoretical foundations for QPINN convergence, proving that quantum advantage persists even under realistic noise conditions characteristic of near-term quantum hardware. Our analysis reveals universal scaling properties: QPINNs consistently achieve mean squared errors smaller than those of classical PINNs, while requiring logarithmically fewer computational resources relative to problem dimensionality. The ubiquity of these improvements across fundamentally different mathematical structures—from hyperbolic conservation laws to elliptic field equations to parabolic diffusion processes—suggests deep connections between quantum computational advantages and the universal mathematical principles underlying diverse physical phenomena.
Authors: Muhammad Al-Zafar Khan, Jamal Al-Karaki, Assala Benmalek, Abdullah Al Omar Ghalib and Marwan Omar

Title: Topological data analysis with variational quantum circuits

Abstract: Understanding the topological structure of high-dimensional datasets—often robust to noise and small-scale deformations—has become increasingly important in modern data science. Topological Data Analysis (TDA) provides a framework for extracting such global features using tools from algebraic topology. A key invariant in TDA is the Betti number, which counts the number of k-dimensional holes in a simplicial complex constructed from a proximity graph on the dataset. However, classical algorithms for computing Betti numbers scale exponentially with either the dimension k or the number of vertices. While several quantum algorithms have been proposed to achieve potential quantum advantage, the problem can also be reformulated as determining the ground-state degeneracy of a data-dependent fermionic Hamiltonian. Motivated by this connection, we explore a variational quantum approach to detect the nontriviality of Betti numbers in clique complexes. Our method employs a parameterized quantum circuit with a particle-number-preserving ansatz to approximate ground states, offering a quantum-inspired alternative for applications in persistent homology.
Authors: Zherui Wang, Jordi Tura and Patrick Emonts

Title: Topological Signal Processing on Quantum Computers for Higher-Order Network Analysis

Abstract: Predicting and analyzing global behaviour of complex systems is challenging due to the intricate nature of their component interactions. Recent work has started modelling complex systems using networks endowed with multiway interactions among nodes, known as higher-order networks. Simplicial complexes are a class of higher-order networks that have received significant attention due to their topological structure and connections to Hodge theory. Topological signal processing (TSP) utilizes these connections to analyze and manipulate signals defined on non-Euclidean domains such as simplicial complexes. Such analysis of higher-order network data is important for many real-world problems, such as detecting failure/error in communication networks, sensor coverage analysis, statistical ranking problems, finding arbitrage currency markets, etc. However, as the dimension of the higher-order networks increases, the complexity of TSP scales exponentially. In this work, we present a general quantum algorithm for implementing filtering processes in TSP and describe its application to extracting network data based on the Hodge decomposition. We leverage pre-existing tools introduced in recent quantum algorithms for topological data analysis and combine them with spectral filtering techniques using the quantum singular value transformation framework. While this paper serves as a proof-of-concept, we obtain a super-polynomial improvement over the best known classical algorithms for TSP filtering processes, modulo some important caveats about encoding and retrieving the data from a

quantum state. The proposed algorithm generalizes the applicability of tools from quantum topological data analysis to novel applications in analyzing high-dimensional complex systems.
Authors: Caesnan Leditto, Angus Southwell, Behnam Tonekaboni, Muhammad Usman, Kavan Modi and Gregory White

Title: Toward Undetectable Backdoors in Variational Quantum Models

Abstract: Setting: Variational Quantum Models (VQMs), such as Variational Quantum Classifiers and Variational Quantum Linear Solvers, represent a promising tool in the emerging field of quantum machine learning and can potentially be used for a large variety of sectors and purposes. However, for the foreseeable future, only a select number of entities and organizations have access to machines that can run them. Therefore, training and inference of such models need to be delegated, outsourced, or externally influenced by those organizations, which may have ulterior motives than those of the original model owner. As such, understanding this vulnerability, particularly in adversarial settings, is of utmost importance. In Goldwasser et al., this threat was made explicit, primarily in classical machine learning. The authors demonstrated the construction of undetectable backdoors, a mechanism that allows an attacker to perturb any input into an adversarial version using a secret key, in order for the model to give a desired output. This is done while the backdoor appears undetectable to any polynomial-time auditor without this key, using cryptographic tools like digital signatures. The authors defined two different types of undetectability of the backdoors. The black-box undetectable backdoor is a backdoor that cannot be detected by a verifier who can query the model's outputs but doesn't know its internal structure. The white-box undetectable backdoor remains undetectable even if the model's full code, parameters, and training data are visible to the verifier. Using these definitions, the authors provided a variety of constructions for both black- and white-box undetectable backdoors for classical machine learning models. In the quantum setting, cryptographic protections and attack vectors are still in their early stages, and there are no known frameworks for embedding cryptographically secure or undetectable control mechanisms inside VQMs. Moreover, most VQMs rely on classical optimizers, exposing them to classical adversarial manipulations that propagate to the quantum layer. Given this, in this work, we ask: How can we implement undetectable backdoors in VQMs? Specifically, we study the construction of undetectable backdoors in VQMs, and under which conditions can we "whiten" a black-box undetectable backdoor. In doing so, we explore the quantum analogues of digital signatures, what level of undetectability can be achieved using each of these constructions, and to which kind of data. Impact and outlook: By adapting classical undetectable backdoor techniques to quantum machine learning, we aim to demonstrate concrete attack opportunities in current QML architectures, and motivate the development of quantum-native model certification mechanisms.
Authors: Eleanor Kedem, Ryan Sweke and Francesco Petruccione

Title: Towards a Framework for Analyzing Quantum Machine Learning Algorithms

Abstract: Quantum machine learning (QML) has recently shown advances along two main tracks. Kernel-based methods, often combined with classical post-processing, have demonstrated problem-specific success and even quantum advantages. In contrast, variational approaches offer broader applicability but often struggle when applied blindly to unfamiliar data. The gap between these paradigms—often navigated via heuristics—remains poorly understood, hindering the development of robust and general-purpose QML algorithms. In this talk, I will present a systematic approach to analyzing and guiding the design of heuristic QML models through numerical diagnostics. I will outline key properties such tools should capture, and showcase three

examples: (i) assessing the average gradient magnitude across optimization landscapes, (ii) testing the randomness of generated quantum states, and (iii) evaluating how data-induced randomness influences classification performance. These tools mark a step toward a general-purpose characterization framework for QML models, aiming to inform both theoretical development and experimental implementation.
Authors: Adrián Pérez-Salinas, Berta Casas, Xavier Bonet-Monroig and Hao Wang

Title: Training parameterized quantum circuits with forward gradients

Abstract: Current gradient-based methods for training quantum circuit-based models for machine learning, optimization, or chemistry tasks face challenges due to the accessibility of gradients from the models. In this work, we introduce a gradient estimation method inspired by forward propagation methods in classical machine learning. Specifically, we aggregate random directional derivatives of quantum circuit parameters to create unbiased gradient estimators, which can lead to improved convergence with fewer measurement shots than alternative estimators. We use these to construct a novel optimizer that adaptively alters the number of measurement shots required during training, enabling convergence with minimal resource requirements. Finally, we show how this method unifies several gradient-based estimators from the literature, including SPSA, parameter-shift, and random coordinate descent methods.
Authors: Brian Coyle, Snehal Raj and El Amine Cherrat

Title: Universal and Efficient Quantum State Verification via Schmidt Decomposition and Mutually Unbiased Bases

Abstract: Efficient verification of multipartite quantum states is crucial to many applications in quantum information processing, but most verification protocols known so far are tailored to quantum states with special structures. By virtue of Schmidt decomposition and mutually unbiased bases, here we propose a universal verification protocol based on adaptive local projective measurements that applies to arbitrary multipartite quantum states. Moreover, we establish a universal upper bound on the sample complexity, which is independent of the local dimensions. Numerical calculations further indicate that the sample complexity for Haar-random pure states is independent of the qudit number. In addition, we propose a simpler verification protocol based on mutually unbiased bases, which avoids Schmidt decomposition and is more amenable to experimental realization, but can achieve a similar efficiency according to numerical simulations. Our work suggests that most multipartite pure states can be verified using a constant sample cost, irrespective of the qudit number and local dimensions.
Authors: Yunting Li and Huangjun Zhu

Title: Universal approximation of continuous functions with minimal quantum circuits

Abstract: The conventional paradigm of quantum computing is discrete: it utilizes discrete sets of gates to realize bitstring-to-bitstring mappings, some of them arguably intractable for classical computers. In parameterized quantum approaches, the input becomes continuous and the output represents real-valued functions. While the universality of discrete quantum computers is well understood, basic questions remain open in the continuous case. We focus on universality on multivariate functions. Current approaches require either a number of qubits scaling linearly with the dimension of the input for fixed encodings, or a tunable encoding procedure in single-qubit circuits. The question of whether universality can be reached with a fixed encoding and sub-linearly many qubits remained open for the last five years. In this paper, we answer this

question in the affirmative for arbitrary multivariate functions. We provide two methods: (i) a single-qubit circuit where each coordinate of the arguments to the function to represent is input independently, and (ii) a multi-qubit approach where all coordinates are uploaded simultaneously, with a number of qubits scaling logarithmically with the dimension of the argument of the function of interest. We view the first result of inherent and fundamental interest, whereas the second result opens the path towards representing functions whose arguments are densely encoded in a unitary operation, possibly encoding for instance quantum processes.
Authors: Adrián Pérez-Salinas, Mahtab Yaghubi Rad, Alice Barthe and Vedran Dunjko

Title: Unsupervised domain adaptation for quantum data under state and distribution shifts

Abstract: An emerging paradigm in practical quantum machine learning is the combination of efficient classical feature extraction from quantum data with classical machine learning techniques. This approach is particularly promising in scenarios where quantum data encode properties that are classically intractable but can be accessed efficiently using quantum hardware. A notable example is the use of classical shadows to extract features from quantum states for classification tasks in many-body physics. However, in realistic applications, obtaining clean and fully labeled quantum data from the complex target domain is often infeasible. This challenges the direct applicability of supervised learning frameworks typically assumed in prior works. To address this issue, we consider an unsupervised domain adaptation setting involving quantum data, where the source domain consists of labeled quantum states and the target domain consists of unlabeled quantum states. Notably, both the data distributions and the quantum states differ between the two domains, introducing state and distribution shifts. We propose to apply domain adaptation techniques to the classical shadows of quantum data.
Authors: Kosuke Ito, Hiroshi Yano, Yudai Suzuki, Akira Tanji and Naoki Yamamoto

Title: Usable Information in Quantum Machine Learning

Abstract: Quantum machine learning (QML) models have been theoretically proven to possess interesting properties such as the potential for enhanced model capacity compared to classical counterparts. These proofs, however, are mainly focused on the properties of a general model class and hence do not incorporate information about specific problems such as the given data or the model sub-space that is accessed during and/or after the training. To gain insights into the practical behaviour of QML models for actual problem settings, one needs to test scalable QML algorithms on non-trivial system sizes with quantum hardware and evaluate suitable metrics. Recent advancements in quantum computing hardware bring us closer to being able to run these tests. Traditional machine learning performance metrics—such as accuracy, precision, and F1 score—focus solely on output predictions and fail to reflect a model's ability to internalise the functional relationship between inputs and outputs. To address this gap, this work introduces and investigates the concept of the V-usable information for QML models. The V-usable information, a metric that quantifies the amount of information a model class can extract from a dataset, is applicable to both classical and quantum models, enabling fair and meaningful comparisons. It extends Shannon's mutual information and can be reliably estimated using the Probably Approximately Correct (PAC) framework, even in high-dimensional settings. This makes it particularly suitable for analysing model expressivity, training dynamics, and generalisation performance. The metric also enables fine-grained analysis of model behaviour, including sensitivity to hardware noise and overfitting. Furthermore, it supports a paradigm shift in generalisation theory—one that emphasises the role of data distribution and trained model behaviour over abstract model class assumptions. As such, the V-usable information emerges as an

interesting tool for advancing our understanding of quantum learning systems and their potential to outperform classical approaches in practical, data-driven scenarios.
Authors: M. Emre Sahin, Cenk Tüysüz, Edoardo Altamura, David Galvin, Oscar Wallis, Alex Gregory, Paul Edwards, Stefano Mensa and Christa Zoufal

Title: Variational bounds of quantum information-theoretic quantities for quantum error correcting code analytics

Abstract: The main advantage of quantum computing over classical information processing is considered to be the increased informational content of quantum states manifested by correlations that go beyond classical correlations—entanglement. Quantifying correlations in quantum systems is an important problem in quantum information processing. In particular, assessing the decay of correlations in a quantum state under environmental noise is crucial for developing and characterizing quantum error-correcting codes and fault-tolerant quantum circuit designs. Such scenarios correspond to information transmission through a noisy channel and can be described using coherent information, which captures the amount of recoverable information after the application of a noisy quantum channel, and quantum mutual information, which quantifies the full amount of correlations in a quantum state. These are entropic quantities whose exact calculation requires computationally expensive quantum state tomography, making them impractical for large-scale systems relevant to logical encodings. Good bounds for coherent and quantum mutual information thus provide useful tools for practical assessment of quantum systems at relevant scale. In classical information theory, significant advances have been made in deriving rigorous variational bounds on mutual information, which can be leveraged in studies of complex physical systems. Similar progress in quantum information theory has so far been limited to variational bounds for quantum relative and entanglement entropies. In this work, we explore multiple variational bounds of quantum entropies and coherent information based on positive operator-valued measurements and purity estimations of quantum states. We apply a hybrid quantum-classical machine learning approach to bound quantum entropies and use the estimated coherent information to derive error thresholds for the rotated surface code and the color code in settings of quantum memory and logical computation under realistic noise models. Leveraging these techniques, we estimate quantum mutual information and apply it to quantify spatial correlations in noisy quantum dynamics, showcasing limitations of quantum error-correcting codes under correlated noise.
Authors: Vladyslav Los and Maciej Koch-Janusz

Title: Variational quantum Hamiltonian engineering

Abstract: The expectation value of a Hamiltonian and Hamiltonian simulation are two fundamental tasks in quantum computation, with their efficiency heavily dependent on the Pauli norm of the Hamiltonian, which sums the absolute values of its Pauli coefficients. In this work, we propose a variational quantum algorithm called Variational Quantum Hamiltonian Engineering (VQHE) to minimize the Pauli norm of a Hamiltonian, thereby reducing the overhead for expectation value estimation and Hamiltonian simulation. We develop a theoretical framework that transforms Pauli norm minimization into a vector l1-norm minimization problem, design an appropriate cost function, and employ parameterized quantum circuits to variationally minimize this cost. Numerical experiments demonstrate the effectiveness of VQHE in reducing the Pauli norm for both the Ising Hamiltonian and molecular Hamiltonians. Furthermore, VQHE is compatible with grouping strategies, enabling further reductions in measurement complexity for expectation value

estimation. Our results highlight the potential of VQHE to enhance the efficiency of quantum algorithms on near-term devices, offering a promising approach for practical applications.
Authors: Benchi Zhao and Keisuke Fujii

Title: Variational quantum self-attention for the prediction of classical and quantum data sequences

Abstract: We propose a quantum implementation of self-attention—the mechanism underlying transformers and large language models—based on a variational quantum circuit. Using this, we construct a variational quantum transformer (VQT) that predicts future data points from input classical or quantum data sequences. Simulations show that the VQT learns small datasets of classical sentences as well as sequences of quantum states evolved under random non-local Pauli Hamiltonians.
Authors: Alessio Pecilli and Matteo Rosati

Title: Variational Quantum Solver for Time-Fractional Differential Equations with Efficient Memory Scaling

Abstract: Fractional differential equations (FDEs) are increasingly used to model complex systems with memory and nonlocal behavior. Classical numerical methods for time-fractional equations face prohibitive memory costs due to the global history dependence of the Caputo derivative. We present a hybrid quantum-classical algorithm that leverages variational quantum optimization to solve FDEs using parameterized quantum circuits. The quantum algorithm shows favorable time complexity and sub-exponential memory scaling compared to classical methods. We demonstrate the approach on benchmark problems, including the fractional diffusion equation, nonlinear Burgers' equation, and a coupled SEIR epidemic model. This work suggests that variational quantum algorithms could offer a route toward solving high-dimensional or memory-intensive FDEs.
Authors: Fong Yew Leong, Dax Enshan Koh, Jian Feng Kong, Siong Thye Goh, Jun Yong Khoo, Wei-Bin Ewe, Hongying Li, Jayne Thompson and Dario Poletti

Title: Weak Form Regularisation for Solving Differential Equations with Quantum Neural Networks

Abstract: Algorithms utilising quantum computers have been proposed to solve differential equations. One such proposition is differentiable quantum circuits: a variational, loss-function minimisation approach, similar to classical physics-informed methods. Such loss functions are local and based on collocation points chosen in the function domain. This emphasises precise solving of differential equations at these points, but in turn may fail at properly propagating boundary conditions, as well as generalising to unseen points. In classical differential equation solvers, a complementary approach entails studying the weak form of the differential problem—i.e., a global, integral-based approach. This encourages generalisation and boundary propagation to ensure a best average fitting over the whole domain, rather than precise fitting in sampled points. Inspired by such complementarity, here we explore combining contributions from both loss-function terms in a unified framework, to reap the advantages whilst mitigating the weaknesses of both. We showcase the robustness of this approach on a variety of problems, in particular adopting domain decomposition strategies.
Authors: Annie Paine, Smit Chaudhary and Antonio Gentile