Funded by
the European Union
NextGenerationEU

Ministero
dell'Università
e della Ricerca

Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA

FAIR
Future
Artificial
Intelligence
Research

Quantum Machine Learning Conference
GGI, Florence

June 25 2025-June 27 2025

	Wednesday 25	**Thursday 26**	**Friday 27**
9.00-9.50	HOLMES	CARO	PASTORELLO
9.50-10.40	ANGRISANI	GIRARDI	SUTTER
10.40-11.10	Break	Break	Break
11.10-12.00	JERBI	OTTAVIANI (remote talk)	MAZZUCCHI
12.00-12.50	ANSCHUETZ	MACALUSO	PEREIRA
12.50-14.50	Lunch	Lunch	Lunch
14.50-15.40	MELE	MARI	
15.40-16.30	STILCK FRANCA	MORGILLO	
16.30-18.00	Coffee & posters	Coffee & posters	

Wednesday 25

Time	**Title**	**Speaker**
9.00-9.50	Pauli and Majorana Propagation methods for classically simulating quantum circuits	HOLMES

Abstract

Simulating quantum circuits classically is in general a hard task. However, certain families of quantum circuits may be practically or even provably efficiently simulable by use of specialized classical algorithms. In this talk, we will cover ”Pauli propagation” which has recently been shown to enable efficient classical simulation of expectation values in quantum circuits and a wide range of noise-free quantum circuits. Appreciating the strengths and weaknesses of this simulation method, and how it can be efficiently combined with other classical and quantum subroutines, will help point towards promising applications of quantum devices. We will end by discussing a generalization of this approach to Fermionic systems opening up new applications in quantum chemistry and material science. This talk will give an overview of the following works: arxiv:2308.09109, arXiv:2408.12739, arXiv:2409.01706, arXiv:2411.19896, arXiv:2501.13101, arXiv:2503.18939.

Time	**Title**	**Speaker**
9.50-10.40	On the interplay between noise, scrambling and classical simulation of quantum systems	ANGRISANI

Abstract

Simulating arbitrary quantum dynamics with classical algorithms is widely believed to be intractable. Yet, by exploiting the structure of certain restricted settings, specialized classical methods can succeed. One particularly promising family-Pauli propagation - recasts simulation in the Pauli basis and often delivers rigorous runtime and error guarantees. At the heart of these guarantees lie two ingredients: the presence of local noise, which dampens long-range interactions, and a degree of "scrambling" in the circuit's gates. In this talk, we will present our recent results on applying Pauli propagation to both noisy and noiseless circuits. Along the way, we'll discuss to what extent noise and scrambling influence simulability - and what that tells us about the necessary resources for quantum advantage.

Time	**Title**	**Speaker**
11.10-12.00	Shadows of quantum machine learning and shallow-depth learning separations	JERBI

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, contingent on widely believed assumptions in complexity theory. In the second work, we expand our understanding of where quantum advantages can be found in quantum machine learning, by showing 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. The machine learning task we consider is that of learning probability distributions, or generative learning. We design this learning task building on recent results by Bene Watts and Parham on quantum advantages for sampling, which we technically uplift to a hyperplane learning problem, identifying non-local correlations as the origin of the quantum advantage.

Time	**Title**	**Speaker**
12.00-12.50	A Unified Theory of Quantum Neural Network Loss Landscapes	ANSCHUETZ

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. While there are settings where quantum neural networks (QNNs) have also been shown to behave as Gaussian processes, there exist known counterexamples to this behavior. 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.

Time	**Title**	**Speaker**
14.50-15.40	Effect of noise in typical quantum circuits	MELE

Abstract

Motivated by realistic hardware considerations of the pre-fault-tolerant era, we comprehensively study the impact of uncorrected noise on typical quantum circuits. We first show that any noise ‘truncates’ most quantum circuits to effectively logarithmic depth, in the task of estimating observable expectation values. We then prove that quantum circuits under any non-unital noise exhibit lack of barren plateaus for cost functions composed of local observables. But, we also design an efficient classical algorithm to estimate observable expectation values of typical quantum circuits within any target constant accuracy, in any circuit architecture. Taken together, our results showcase that, unless we carefully engineer the circuits to take advantage of the noise, it is unlikely that noisy quantum circuits provide any quantum advantage for algorithms that output observable expectation value estimates, like many variational quantum machine learning proposals.

Time	**Title**	**Speaker**
15.40-16.30	Information-theoretic Generalization Bounds for Learning from Quantum Data	STILCK FRANCA

Abstract

Learning tasks are playing an increasingly central role in quantum information and computation—from fundamental problems like state discrimination and metrology to quantum PAC learning and the recently proposed “shadow” variants of state tomography. Yet these various strands of quantum learning theory have largely evolved in isolation. In this talk, we introduce a unified mathematical framework for quantum learning based on classical–quantum training data and show how to bound a quantum learner’s expected generalization error on new data. Our bounds are expressed in terms of classical and quantum information theoretic quantities that capture how strongly the learner’s hypothesis depends on the specific training data. To derive them, we develop non commutative analogues of the decoupling lemmas underlying recent classical information theoretic generalization bounds, drawing on tools from quantum optimal transport and quantum concentration inequalities. This framework subsumes and yields intuitive generalization bounds for a variety of quantum learning scenarios—including quantum state discrimination, PAC learning of quantum states or classical functions, and quantum parameter estimation—laying the groundwork for a unified, information theoretic perspective on quantum learning.

Time	
16.30-18.00	**Coffee & posters**

Thursday 26

Time	**Title**	**Speaker**
9.00-9.50	Interactive proofs for verifying (quantum) learning and testing	CARO

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.

Time	**Title**	**Speaker**
9.50-10.40	Trained quantum neural networks and Gaussian processes	GIRARDI

Abstract

We study quantum neural networks made by parametric one-qubit gates and fixed two-qubit gates in the limit of infinite width, where the generated function is the expectation value of the sum of single-qubit observables over all the qubits. First, we prove that the probability distribution of the function generated by the untrained network with randomly initialized parameters converges to a Gaussian process whenever each measured qubit is correlated only with few other measured qubits. Then, we analytically characterize the training of the network via gradient descent with square loss on supervised learning problems. In particular, as long as the network is not affected by barren plateaus, the trained network can perfectly fit the training set and that the probability distribution of the function generated after training still converges in distribution to a Gaussian process, also in the presence of the statistical noise of the measurement at the output of the network. For finite size circuits, we make the convergence quantitative in terms of the Wasserstein distance of order 1.

Time	**Title**	**Speaker**
11.10-12.00	EuroQHPC-I and QEC4QEA: role of Italy into the european HPC-QC ecosystem	OTTAVIANI

Abstract

Italy is at the forefront of shaping the European HPC-QC ecosystem, playing a key role in two major initiatives: EuroQHPC-I and QEC4QEA. As one of the selected hosting entities for a European quantum computer, Italy is set to pioneer the integration of quantum computing with high-performance computing (HPC). This integration, conducted alongside other selected hosting entities, will mark a significant step toward the hybrid computing architectures of the future. Simultaneously, Italy has been chosen to lead Europe's first Center of Excellence in Quantum Computing, QEC4QEA. This initiative will drive the development of the first HPC-QC applications, accelerating the adoption of quantum technologies in scientific and industrial domains. By spearheading both infrastructure deployment and software innovation, Italy is in pole position to build the future European HPC-QC ecosystem, reinforcing its leadership in quantum and high-performance computing.

Time	**Title**	**Speaker**
12.00-12.50	Limitations and Methodological Pathways in Quantum Unsupervised and Reinforcement Learning	MACALUSO

Abstract

Quantum Machine Learning (QML) has recently been explored as a novel approach to surpass the capabilities of classical methods, although the field remains in its early stages and the outcomes achieved so far are still inconclusive. This talk offers a critical and methodologically grounded perspective on current QML approaches, with particular attention to the fundamental limitations of classical machine learning and the ways in which quantum-enhanced models may be designed to address these challenges. Recent developments in unsupervised and reinforcement learning serve as illustrative examples to examine how quantum formulations, tailored to the structure of specific problems, can yield algorithmic and representational advantages. Methodological aspects such as model design, problem encoding, and hybrid integration are emphasized, along with a discussion of current limitations in quantum computing, including hardware constraints and the lack of mature, task-specific quantum design principles. The talk concludes with reflections on how these insights may inform the development of more robust and effective QML methodologies.

Time	Title	Speaker
14.50-15.40	Can AI learn the best way to use a noisy quantum computer?	MARI

Abstract

We explore the broad question posed in the title from different perspectives. We show how a classical neural network can be trained to optimally embed features into a quantum system and to optimally extract information from it. We review the concept of variational quantum error mitigation, i.e., the idea of variationally optimizing error mitigation strategies. We present recent results demonstrating how classical deep learning models and noisy quantum computers can cooperate to better estimate quantum expectation values. Finally, as a speculative open problem, we propose pushing the core question to its extreme limit: Can AI autonomously decide how to optimally use a noisy quantum computer without hard-coding any specific error-reduction strategy?

Time	Title	Speaker
15.40-16.30	Quantum optical classifier with superexponential speedup	MORGILLO

Abstract

Cat or dog? Can a Hong-Ou-Mandel interferometer tell the difference? This talk presents a quantum optical method for binary classification that recognizes patterns without the need for image reconstruction. By encoding both data and model parameters into single-photon states and leveraging two-photon interference, the system classifies patterns directly through coincidence rates. Acting as a quantum analogue of a classical neuron, it operates—once trained—with constant O(1) resource complexity, achieving a superexponential speedup over its classical counterpart.

Time	
16.30-18.00	**Coffee & posters**

Friday 27

Time	Title	Speaker
9.00-9.50	Quantum GANs and optimal mass transport	PASTORELLO

Abstract

After an introduction to the notion of quantum generative adversarial networks (qGANs), I will summarize a recent quantum tomography protocol for constructing a classical estimate of a quantum state by performing repeated measurements on a n-qubit system. I will then discuss the convergence of the protocol with respect to a quantum version of the first-order Wasserstein distance, a fundamental notion of the theory of optimal mass transport. In particular, I will show how this convergence result allows us to conclude that a qGAN can be equivalently trained using classical estimators of quantum states instead of quantum data. This fact is important in practice, as it enables the training of quantum models without requiring direct access to quantum memory or coherent quantum data streams.

Time	**Title**	**Speaker**
9.50-10.40	Uhlmann's theorem for relative entropies	SUTTER

Abstract

Uhlmann's theorem states that, for any two quantum states ρAB and σA, there exists an extension σAB of σA such that the fidelity between ρAB and σAB equals the fidelity between their reduced states ρA and σA. In this work, we generalize Uhlmann's theorem to α-Rényi relative entropies for $\alpha \in [12, \infty]$, a family of divergences that encompasses fidelity, relative entropy, and max-relative entropy corresponding to $\alpha = 12$, $\alpha = 1$, and $\alpha = \infty$, respectively.

Joint work with Giulia Mazzola and Renato Renner.

Time	**Title**	**Speaker**
11.10-12.00	A swap test photonic integrated circuit for determining quantum entanglement	MAZZUCCHI

Abstract

Entanglement is a fundamental resource in quantum computation and quantum communication, but it is potentially affected by decoherence phenomena that make it necessary to introduce appropriate tests to certify and quantify the degree of entanglement of a quantum state. In this talk. I will show how a photonic integrated circuit designed to implement the swap test algorithm can be adapted to an efficient entanglement witness for both pure and mixed bipartite states.

Time	**Title**	**Speaker**
12.00-12.50	Out-of-distribution generalisation for learning quantum channels with low-energy coherent states	PEREIRA

Abstract

Investigating the input-output relations of a quantum process can be seen as a learning problem. For instance, we could wish to find the optimal parameters for some quantum device that let it best mimic our target process, or we could simply wish to construct the best possible mathematical model of the process. Experimentally, we send probes through the quantum channel and use the outputs as our training set. When learning the action of a continuous variable (CV) quantum process in this way, there will often be some restriction on the input states used. One experimentally simple way to probe CV channels is using low-energy coherent states. Learning a quantum channel in this way presents difficulties, since two channels may act similarly on low energy inputs but very differently for high energy inputs. They may also act similarly on coherent state inputs but differently on non-classical inputs. Extrapolating the behaviour of a channel for more general input states from its action on the far more limited set of low energy coherent states is a case of out-of-distribution generalisation. To be sure that such generalisation gives meaningful results, one needs to relate error bounds for the training set to bounds that are valid for all inputs. We show that for any pair of channels that act sufficiently similarly on low energy coherent state inputs, one can bound how different the input-output relations are for any (high energy or highly non-classical) input. This proves out-of-distribution generalisation is always possible for learning quantum channels using low energy coherent states.