

Reinforcement learning-based quantum feature map search leveraging magic and generalization bounds

Jan Schnabel¹, João Bravo², Ernesto Palidda³, David Garvin³, Marco Roth², and Marco Painsi³

¹Fraunhofer Institute for Manufacturing Engineering and Automation IPA

²Fraunhofer Institute for Industrial Engineering IAO

³Rigetti Computing

Identifying genuine quantum resources that lead to quantum transformations which simultaneously show a separation between classical and quantum learners and cannot be efficiently dequantized is a fundamental challenge in quantum machine learning (QML). Superposition and entanglement alone do not suffice. To characterize the expressive power of a QML model, one must instead consider the complexity of the quantum states it can generate. This is formalized by magic – a quantity that measures the hardness of classical simulation of quantum states. Quantum feature maps with higher magic are those that are hard to simulate classically. Thus, there may be a clear gap in the expressive power of the quantum and classical models, unless classical models are given exponentially larger resources.

Building on the work in Refs. [1,2], magic may serve as a valuable resource in QML: the achievable tradeoff between generalization capacity and approximation capacity is controlled by the magic of the underlying quantum feature map. Importantly, this framework provides a formal, model-independent criterion for distinguishing between promising and unpromising quantum feature maps, decoupled from any specific downstream QML model. However, systematically designing data encoding circuits for a given dataset that are both high in magic and generalize well for unseen data points remains a non-trivial challenge.

In this work, we propose a MuZero-based [3] reinforcement learning (RL) approach, following the architecture search framework introduced in Ref. [4], to address this problem. For this, we further draw on extended theoretical insights from [1] in terms of approximation functionals, generalization factors and a conjectured bound for the product of approximation and generalization. By incorporating both magic and generalization bounds directly into the reward signal, we demonstrate that the resulting RL agent can generate high-magic encoding circuits with favorable generalization properties, without any bias toward a specific downstream QML model. We examine different reward designs that vary the relative weighting of the two competing objectives and show that it is possible to jointly optimize for both. We benchmark the generated circuits across selected datasets and multiple QML models, showing that some models are better suited to particular encoding strategies than others, and we provide explicit comparisons against classical baselines.

Our results suggest that RL-based architecture search provides a systematic, dataset-specific approach to encoding circuit design and a practical tool for probing quantum-classical separations in (un-)supervised learning tasks.

References:

- [1] M. Painsi (2025, March, 21). Magic in Quantum Machine Learning [Oral presentation]. APS Global Physics Summit 2025, Anaheim, CA
- [2] D. Garvin, M. Fiorentini, O. Kondratyev, M. Painsi. Data Anonymization with the Density Matrix Classifier. *The Journal of FinTech* **5** (2), 2550010 (2025)
- [3] Schrittwieser, J., Antonoglou, I., Hubert, T. et al. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature* **588**, 604–609 (2020)
- [4] F. Rapp et al. Reinforcement learning-based architecture search for quantum machine learning. *Mach. Learn.: Sci. Technol.* **6** 015041 (2025)