

How Much Does Nonlinearity Cost? Resource-Aware Quantum Kolmogorov–Arnold Networks

Petr Ivashkov,^{1,2} Po-Wei Huang,^{1,3} Kelvin Koor,¹ Lirandë Pira,¹ and Patrick Reberstrost^{1,4}

¹Centre for Quantum Technologies, National University of Singapore, Singapore

²Department of Information Technology and Electrical Engineering, ETH Zürich, Zürich, Switzerland

³Mathematical Institute, University of Oxford, Oxford, United Kingdom

⁴Department of Computer Science, National University of Singapore, Singapore

How much does nonlinearity cost on a quantum computer? For hardware developers, nonlinear algorithms translate into deeper circuits, more ancillas, and higher error-correction overhead. This directly impacts cooling power and system-level energy consumption. For algorithm designers, nonlinear transformations are essential for expressive quantum machine learning and data processing, but are often proposed without transparent resource accounting. A principled connection between algorithmic design choices and their physical resource costs is central to energy-efficiency in quantum computing.

Design of nonlinearity in quantum learning. We introduce Quantum Kolmogorov–Arnold Networks (QKAN) [1], a fault-tolerant framework for implementing programmable nonlinear transformations on quantum data using block-encodings and Quantum Singular Value Transformation (QSVT) [2]. Inspired by Kolmogorov–Arnold Networks in classical machine learning [3], QKAN decomposes multivariate functions into sums of trainable univariate nonlinearities [4]. This structure is both expressive and resource-structured: each nonlinearity can be implemented independently, enabling control over circuit depth and ancilla usage. As a result, QKAN provides a concrete way to explore energy and accuracy trade-offs in large-scale quantum algorithms and quantum machine learning workflows.

From the quantum software stack perspective, QKAN is built entirely from standard fault-tolerant primitives (block-encodings, QSVT, linear combinations of unitaries), making it compatible with compilation pipelines and error-corrected architectures. This positions QKAN as a reusable nonlinear primitive for hybrid HPC–quantum workflows, where classical preprocessing and quantum subroutines must be co-designed under energy and resource budgets. Beyond learning tasks, QKAN can be used as a programmable nonlinear routine for loading data and transforming features for quantum algorithms, enabling multivariate amplitude transformations that are otherwise difficult to realize.

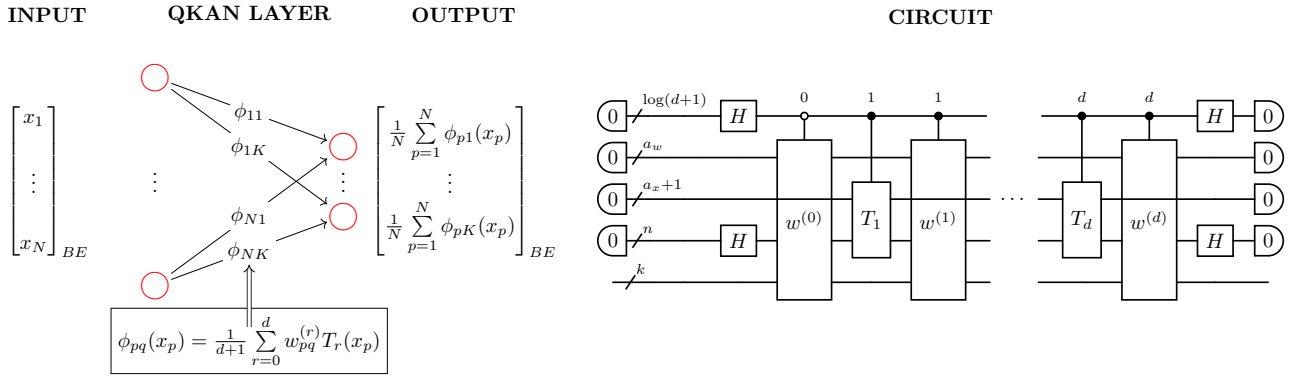


FIG. 1. Construction of a model layer with the corresponding quantum circuit.

Cost of non-linearity. We construct an L -layer QKAN in which each activation function is approximated by a Chebyshev polynomial of degree d implemented via QSVT, shown in Fig. 1. The resulting block-encoding requires $\mathcal{O}\left(\frac{d^{2L}}{2^L}\right)$ controlled uses of input and weight block-encodings, with ancilla overhead linear in L and logarithmic in the input and output dimension. Classical readout via amplitude estimation scales as $\mathcal{O}(1/\delta)$ for additive error δ , while quantum output states can be prepared for downstream quantum subroutines. These explicit scalings allow direct translation from model expressivity to fault-tolerant circuit depth and ancilla costs, enabling back-of-the-envelope estimates of physical resource and energy overheads.

-
- [1] P. Ivashkov, P.-W. Huang, K. Koor, L. Pira, and P. Rebentrost, [Qkan: Quantum kolmogorov-arnold networks](#) (2024), [arXiv:2410.04435 \[quant-ph\]](#).
 - [2] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 193–204.
 - [3] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljagic, T. Y. Hou, and M. Tegmark, KAN: Kolmogorov-Arnold networks, in *The Thirteenth International Conference on Learning Representations* (2025).
 - [4] A. G. Rattew and P. Rebentrost, Non-linear transformations of quantum amplitudes: Exponential improvement, generalization, and applications (2023), [arXiv:2309.09839 \[quant-ph\]](#).