

Introducing Reduced-Width QNNs, an AI-Inspired Ansatz Design Pattern

Jonas Stein^{1,2,*}^a, Tobias Rohe^{1,*}, Francesco Nappi¹, Julian Hager¹, David Bucher²^b,
Maximilian Zorn¹^c, Michael Kölle¹^d and Claudia Linnhoff-Popien¹^e

¹Mobile and Distributed Systems Group, LMU Munich, Germany

²Aqarios GmbH, Germany

fi

Keywords: Quantum Computing, Variational Quantum Circuits, Circuit Optimization, Variational Quantum Eigensolver.

Abstract: Variational Quantum Algorithms are one of the most promising candidates to yield the first industrially relevant quantum advantage. Being capable of arbitrary function approximation, they are often referred to as Quantum Neural Networks (QNNs) when being used in analog settings as classical Artificial Neural Networks (ANNs). Similar to the early stages of classical machine learning, known schemes for efficient architectures of these networks are scarce. Exploring beyond existing design patterns, we propose a *reduced-width* circuit ansatz design, which is motivated by recent results gained in the analysis of dropout regularization in QNNs. More precisely, this exploits the insight, that the gates of overparameterized QNNs can be pruned substantially until their expressibility decreases. The results of our case study show, that the proposed design pattern can significantly reduce training time while maintaining the same result quality as the standard "full-width" design in the presence of noise. We thus argue, that quantum architecture search should not blindly follow the classical overparameterization trend.


1 INTRODUCTION


A central goal of current quantum computing research is achieving an industrially relevant quantum advantage. As immense computational requirements of known quantum algorithms enabling provable speedups exceed prevailing hardware capabilities in the present NISQ era, less demanding approaches like Quantum Neural Networks (QNNs) are of great interest, especially as they have already shown very promising results towards a possible quantum advantage (Abbas et al., 2021). QNNs closely resemble classical Artificial Neural Networks (ANNs), as they are composed of concatenated, parameterized functions capable of approximating arbitrary functions (Mitarai et al., 2018). The core difference between QNNs and ANNs is the dimensionality of the underlying vector space, which is exponentially larger


in the former with respect to the space requirements in qubits/neurons (Mitarai et al., 2018).


A key advantage of QNNs compared to established quantum algorithms like those of (Shor, 1994) or (Grover, 1996) is their flexibility towards respecting restricted hardware capabilities. More precisely, QNNs can be shaped to fit a maximal width and depth of the circuit to be executed on the QPU and can thus minimize errors caused, e.g., by decoherence. Analogous to the historical progress on ANNs, where many useful and performant network architectures and design patterns emerged, recent literature on QNNs explores productive quantum circuit design patterns (Cerezo et al., 2021a; Sim et al., 2019a; Schuld et al., 2021; Havlíček et al., 2019). Similar to ANNs, these techniques are used to find a sufficiently good trade-off between expressibility and trainability.


In classical machine learning, this problem is predominantly solved by utilizing heavily overparameterized models and *regularization* techniques to facilitate efficient trainability. A particularly powerful regularization technique is dropout, in which parts of the model are randomly deleted for each training step. Among other effects, this introduces noise to the network architecture and therefore hinders overfit-

^a <https://orcid.org/0000-0001-5727-9151>

^b <https://orcid.org/0009-0002-0764-9606>

^c <https://orcid.org/0009-0006-2750-7495>

^d <https://orcid.org/0000-0002-8472-9944>

^e <https://orcid.org/0000-0001-6284-9286>

*These authors contributed equally.

ting (Srivastava et al., 2014). Meanwhile, even though research on regularization techniques for QNNs is still in its early stages, dropout has also shown first promising results (Scala et al., 2023; Kobayashi et al., 2022). Analyses on instances of overparameterized QNNs in (Scala et al., 2023) show no reduction of expressibility in the presence of substantial gate pruning. In this article, we use these insights as inspiration for an ansatz design pattern, that directly employs a gate-pruned version of overparameterized QNNs. This is explicitly not targeted towards substituting dropout, but meant as an exploration of the necessity of the standard full-width layer design in the first place.

To test the effectiveness of our approach, we evaluate QNNs of a fixed circuit architecture before and after pruning. Our main contributions in this paper can thus be summarized as the following:

- We propose a novel, AI-inspired QNN design pattern, that can notably also be applied as a pruning technique for existent QNNs.
- We conduct an empirical case study of this design pattern for the maximum cut optimization problem in a noisy environment.

2 BACKGROUND

In the following, we describe the core algorithms and techniques employed in our methodology. We show how QNNs can be utilized to solve specific optimization problems like the maximum cut problem as part of a Variational Quantum Eigensolver (VQE).

2.1 Quantum Neural Networks

In current literature, the term Quantum Neural Network is predominantly used to describe a Parameterized Quantum Circuit (PQC) utilized for tasks typical for classical Artificial Neural Networks. Although the terminology suggests a close similarity between QNNs and ANNs, there exist significant differences. A core distinction is that QNNs do not consist of individual neurons, instead they are comprised of qubits, whose states can be changed using (parameterized) quantum gates, that can involve interactions with other qubits. Despite the substantial differences between ANNs and QNNs, there also exist strong commonalities, foremost the shared capability of universal function approximation, i.e., ANNs and QNNs can both be described as parameterized functions, that can approximate arbitrary functions under certain conditions (Schuld et al., 2021). Moreover, by representing QNNs and feedforward ANNs




Figure 1: Example network architecture of the encoder in an autoencoder.

as mathematical functions, we can observe clear similarities. A feedforward neural network is mathematically represented by concatenations of typically non-linear activation functions for all neurons in each layer (Schmidhuber, 2015). In contrast, QNNs typically model concatenations of linear layers with a single non-linear layer at the end, which is responsible for the measurement, i.e., generating a classical output (Mitarai et al., 2018). For examples of feedforward ANNs and QNNs, see Figs. 1 and 2.

Mathematically, the final state of a QNN composed of $n \in \mathbb{N}$ qubits can be formalized as denoted in Eq. 1¹. In this formalization, we started in the standard initial state $|0\rangle^{\otimes n}$ and subsequently applied a parameterized unitary operation $U(\theta)$, with parameters $\theta \in \mathbb{R}^m$, where $m \in \mathbb{N}$ denotes the number of parameters. In practice, U is predominantly comprised of parameterized single-qubit and two-qubit gates to match native gate sets of existing quantum computers.

$$|\psi(\theta)\rangle := U(\theta)|0\rangle^{\otimes n} \quad (1)$$

To extract classical information from the final state of the QNN, we typically measure each qubit $|\psi_i\rangle$ in the computational basis $\{|0\rangle, |1\rangle\}$ using the measurement operator $M = I^{i-1} \otimes \sigma_z \otimes I^{n-i}$, where σ_z denotes the Pauli-Z operator. These measurements yield a probability distribution over bitstrings $|x\rangle \in \mathbb{C}^{2^n}$ representing the output of the quantum circuit, where $|x\rangle$ occurs at probability $|\langle x | \psi(\theta) \rangle|^2$.

Having presented an overview of the circuit components of a QNN, we now address parameter training approaches. Analogous to ANNs, two main types of parameter optimization approaches exist: gradient-based and derivative-free methods (Mitarai et al., 2018). While QNNs can be trained with essentially the same techniques as ANNs, i.e., using approaches like gradient descent, genetic algorithms, and others, known gradient computations for QNNs are fundamentally more expensive. Using backpropagation, the gradient in feedforward ANNs can be calculated as

¹This formalization explicitly disregards intermediate measurements, in-line with predominant literature.

fast as a forward pass of the ANN, whereas all known quantum gradient computations need $O(m)$ executions of the QNN, where $m \in \mathbb{N}$ denotes the number of parameters (Mitarai et al., 2018; Schuld et al., 2019). The standard method to calculate the gradient of a parameterized quantum circuit is the parameter shift rule, which is used to calculate each partial derivative $\partial/\partial\theta_i U(\theta)|0\rangle^{\otimes n}$ based on Eq. 2, given that all parameters are enclosed in single qubit gates and $\theta^\pm := (\theta_1, \dots, \theta_{i-1}, \theta_i \pm \pi/2, \theta_{i+1}, \dots, \theta_m)$ (Mitarai et al., 2018; Schuld et al., 2019).

$$\frac{\partial}{\partial\theta_i} U(\theta)|0\rangle^{\otimes n} = \frac{U(\theta^+)|0\rangle^{\otimes n} - U(\theta^-)|0\rangle^{\otimes n}}{2} \quad (2)$$

Given this linear runtime increase in the gradient computation for QNNs, achieving better solution quality or a quick convergence to sufficiently good parameters is essential for QNNs.

2.2 The Maximum Cut Problem in Quantum Optimization

The *maximum cut* problem, is an NP-hard combinatorial optimization problem with the objective of partitioning the nodes $v_i \in V$ of a given graph $G(V, E)$ into two complementary sets $A \cup B = V$ such that number of edges between A and B is maximal. Encoding a partitioning via a binary vector $s \in \{-1, 1\}^n$, where $n := |V|$ and $s_i = 1 \Leftrightarrow v_i \in A$ and $s_i = -1 \Leftrightarrow v_i \in B$, the search for the optimal partition can be formulated as shown in Eq. 3 (Lucas, 2014). This formulation naturally evolves after observing that $(1-s_i s_j)/2 \in \{0, 1\}$ equals one iff the edge $(v_i, v_j) \in E$ connects A and B .

$$\arg \min \sum_{s \in \{-1, 1\}^n} \sum_{(v_i, v_j) \in E} s_i s_j \quad (3)$$

We have chosen this ± 1 encoding of the solution, as it can directly be transformed into a quantum Hamiltonian \hat{H} that represents the solution space of this optimization problem, which is the standard approach in quantum optimization exploiting that Ising Hamiltonians are isomorphic to the NP-hard QUBO. For this, we merely need to substitute s_i using the map $s_i \mapsto \sigma_i^z := I^{\otimes i-1} \otimes \sigma_z \otimes I^{\otimes n-i}$, where I denotes the two dimensional identity matrix. Therefore, we can define $\hat{H} := \sum_{(v_i, v_j) \in E} \sigma_i^z \sigma_j^z$.

Switching to a 0/1 encoding, i.e., $x \in \{0, 1\}^n$ with $x_i = 0 \Leftrightarrow v_i \in A$ and $x_i = 1 \Leftrightarrow v_i \in B$, \hat{H} can be understood as a mapping between solutions x and the value of the objective function. More specifically, its eigenvalues for a given eigenvector $|x\rangle$ are the values of the objective function, where x represents the 0/1 encoding of the solution.

2.3 The Variational Quantum Eigensolver

The Variational Quantum Eigensolver (VQE) is an algorithm targeted towards finding the ground state of a given Hamiltonian \hat{H} using the *variational method* (Peruzzo et al., 2014). The term *variational method* describes the process of making small changes to a function f to find minima or maxima of a given function g (Lanczos, 2012). In the case of the VQE, $f: \theta \mapsto U(\theta)|0\rangle^{\otimes n}$ describes a parameterized quantum circuit yielding the state $|\psi(\theta)\rangle := U(\theta)|0\rangle^{\otimes n}$ which represents a guess of the arg min of the function $g: |\phi\rangle \mapsto \langle \phi | \hat{H} | \phi \rangle$. More specifically, the VQE uses parameter training techniques on the given ansatz $U(\theta)$ to generate an output state $|\psi(\theta)\rangle$ that approximates the ground state of \hat{H} .

For our purposes, i.e., solving an optimization problem, we need to (1) define a Hamiltonian encoding our objective function and (2) select a suitable ansatz in order to use the VQE algorithm. While we can use the procedure exemplified in Sec. 2.2 to define a Hamiltonian \hat{H} , the selection of the ansatz and thus a circuit architecture is not thoroughly understood yet. Although literature on automatized quantum architecture search tools is emerging (Du et al., 2022), most ansätze are designed based on empiric experience (Sim et al., 2019b).

3 RELATED WORK

A suitable circuit design in QNNs is essential for achieving good results in variational quantum computing and is actively researched in literature (Sim et al., 2019b). While some design patterns like a layerwise circuit structure with alternating single qubit rotations and entanglement gates, have become standard practice (Sim et al., 2019b), the list of established circuit elements still is considerably smaller than that in classical neural networks. A central problem when trying to design ansätze is the trade-off between expressibility and trainability. Reasonably complex models demand for expressive circuits, which in turn can quickly lead to the problem of *barren plateaus*, a trainability issue where the probability that the optimizer-gradient becomes exponentially small is directly dependent on the number of qubits (McClellan et al., 2018). Moreover, barren plateaus are found to be dependent on the choice of the cost function itself (Cerezo et al., 2021b) as well as the number of parameters to optimize simultaneously (Skolik et al., 2021).

Classical deep learning faced a similar problem

in *vanishing gradients*, e.g., for recurrent neural networks (Hochreiter, 1998), which eventually lead to the adoption of more specialized network architectures like the long short-term memory (LSTM) networks (Hochreiter and Schmidhuber, 1997) for sequence processing. Similarly, the current trend of improving variational architectures focuses on optimizing the designs of quantum circuit ansätze, e.g., via optimization procedures starting from randomized circuits (Fösel et al., 2021). In this context, our work shows that pruning gates of the standard full-width circuit architecture can substantially increase trainability while retaining the same solution quality.

4 METHODOLOGY

In the subsequent sections, we describe our reduced-width ansatz design pattern in two different variants as pruned versions of a baseline ansatz. Concluding our methodology, we describe our approach to parameter training and problem instance generation.

4.1 Full-Width Design

In the following, we introduce the design of our baseline ansatz, which we coin *full-width design*. We choose a very simple circuit architecture consisting of `BasicEntanglerLayers`², as it resembles the most basic circuit architecture and is ubiquitously used for constructing variational quantum circuits in quantum machine learning.

More specifically, this baseline circuit is constructed of two components. The first is a sequence of parameterized gates, i.e., a Pauli- X , a Pauli- Z and another Pauli- X rotation, applied to each qubit, allowing for the implementation of arbitrary single qubit operations. The second is a circular CNOT-entanglement, i.e., CNOTs applied to all neighboring qubit pairs, using the maximum width of the circuit. We combine these single- and two-qubit gate components into a single layer, where every circuit is composed of four such layers to analyse scaling behavior. The entire circuit with all its layers is illustrated in Fig. 2. In the following sections, we will refer to this circuit design as the *full-width circuit* or *full-width design*, as it uses the same maximum width throughout all layers. All the circuits utilized in this study are derived from this full-width circuit, where variations are introduced through the removal of gates.

²For details and the implementation, see <https://docs.pennylane.ai/en/stable/code/api/pennylane.BasicEntanglerLayers.html>.

4.2 Reduced-Width Design

In this section, we introduce two different approaches that reduce the width of a given ansatz by pruning gates in a layerwise manner. The first approach works analog to the dropout employed in (Scala et al., 2023) and randomly prunes the a fixed number of gates from the circuit (for reference, see Fig. 3). To test how the structure of the employed pruning affects the performance, we propose a second approach, which gradually increases the number of pruned gates for subsequent layers (for reference, see Fig. 4).

Starting with the detailed discussion of the second approach, we review a key result of (Scala et al., 2023), which states that the deeper the circuits, the more gates can be pruned before effectively reducing the expressibility. This is very sensible, as their results on the effectiveness of pruning are conditioned on present overparameterization, which only occurs for reasonably deep QNNs. Furthermore, we argue, when entanglement between the qubits is sufficiently high (i.e., in general after some ansatz layers), every gate executed on a subset of these qubits intrinsically also acts on all qubits, as they must be regarded as one system. A well-known example in which exactly this happens, is the ancilla quantum encoding as part of the quantum algorithm for solving linear systems of equations (Harrow et al., 2009; Stougiannidis et al., 2023), in which controlled rotations on a single ancilla qubit substantially change the state of an entire qubit register.

In our implementation, we reduce the circuit width in every layer by a magnitude of two in a symmetric manner. This implies that the outermost qubits manipulated in the previous layer are no longer altered directly in the appended layer. This reduction applies to both single- and two-qubit gates. Consequently, the circular entanglement is applied solely to the remaining inner qubits, so that gates are not only deleted, but also reconnected. A visualization of our reducing-width ansatz can be found in Fig. 4.

The number of reduced single- and two-qubits gates is dependent on the number of qubits n , the number of layers l , and the magnitude by which the circuit width is reduced per layer m . Further, we have to distinguish between the reduction in one-qubit gates $\sum_{i=1}^l (i-1)3m$, and the reduction in two-qubit gates $\sum_{i=1}^l (i-1)m$. The proposed reduction is naturally restricted by the number of qubits in the circuit, i.e., $(l-1)m < n$ must be satisfied, so that the reduction can be conducted when starting from the full-width circuit.

In our case (i.e., using eight qubits, four ansatz layers, and a reduction of the circuit width per layer




Figure 2: The full-width design in which gates are applied over the full width of the circuit throughout all layers. The circuit is displayed with four layers, in-line with its configuration in our evaluation.




Figure 3: The random-width design in which a specific number of gates are deleted from the full-width design at random positions. The number of gates removed equals the number of removed gates in the reducing-width circuit. The circuit is displayed with four layers, in-line with its configuration in our evaluation.




Figure 4: The reducing-width design in which the width is reduced from layer to layer by a magnitude of two. With this design approach, we prune the circuit to a version which uses much fewer gates. The circuit is displayed with four layers, in-line with its configuration in our evaluation.

of two in magnitude), we reduce the number of one-qubit gates by 36, and the number of two-qubit gates by 12. Respectively, this corresponds to a reduction of one-qubit gates by 37.5%, and a 50% reduction for two-qubit gates. Beside the hypothesized positive effects that a reduction of the number of gates has on the parameter optimization process, the proposed circuit design is naturally also more noise tolerant. More specifically, the reduced number of gates and consequently the lower circuit depth reduce the total noise accumulated from imperfect gate fidelities and decoherence (Hashim et al., 2023). We therefore expect our approach to be the most beneficial in a noisy environment.

As stated in the beginning of this section, we also introduce the so called random-width design. Here, we prune the same number and the same kind of gates from the full-width design as in the reducing-width approach. However, instead of removing these gates in a systematic way, as done in the reducing-width design, we randomly determine the gates for removal. Therefore, the number and kind of gates compared to the reducing-width design stays the same, but the sequence and position of the gates differ. An exemplary circuit resulting from the random-width design is shown in Fig. 3. To achieve statistical relevance in

the structure of the pruned gates, the results of three different prunings are averaged over each problem instance in the evaluation.

4.3 Training Method

Aiming for a test case in which our approach can show its full potential (see Sec. 4.2), we decide to restrict our evaluation to a noisy environment. Therefore, we employ layerwise learning, which is a method specifically targeted at coping with noise during the training and has been shown to increase result quality significantly in many cases (Skolik et al., 2021). A small prestudy also shows that this is the case for our setup. The training conceptually starts with the first layer of the respective circuit individually, i.e., without any other layers present in the circuit. Subsequently, the parameterized gates get initialized with zeroes to approximate an identity layer similar to the approach proposed in (Liu et al., 2022) and the parameter training is started. For the training process, we use the COBYLA optimizer (Powell, 1994), where the maximum number of iterations is set to 300 to allow for the learning curve to converge sufficiently, according to empirical pretests conducted for this study. After the training of the first

layer is completed, the trained parameters get saved, and a new quantum circuit is constructed. The new circuit now implements two layers and re-uses the previously trained and saved parameters for the initialization of the first layer. The parameters of the gates added in the next layer of the circuit are initialized with zero, while the subsequent training-process only trains the newly added layer of the respective circuit. This training method now repeats itself until all four layers are implemented and trained. This training approach therefore aims at improving on every previously achieved solution iteratively, similar to techniques like reverse annealing.

4.4 Problem Instance Generation

To evaluate the performance of the three different design approaches, we solve the maximum cut problem for 50 randomly generated 8-node Erdős-Rényi graphs (Erdős and Rényi, 1959) using a VQE setup with the respective circuit design as the ansatz. These graphs are generated in a random manner, where the number of nodes, as well as the probability that an edge exists between each pair of nodes, are given as input. We use eight nodes to match the qubit count and choose an edge probability of 50% to yield reasonably realistic problem instances, analogous to (Martiel et al., 2021).

5 EVALUATION

Our experiments are conducted with the three different ansätze shown in Figs. 2, 3 and 4 aiming to solve 50 maximum cut problem instances using a VQE approach and layerwise learning. As motivated in Sec. 4.2, we employ depolarizing noise effects in our circuit simulations: The one-qubit and two-qubit gate error rate were both set to 1%, which roughly coincides with the gate fidelities of recent IBM QPUs. To facilitate these comparatively compute intensive circuit simulations, we run our experiments on an Atos Quantum Learning Machine (QLM).

In the following two sections, we investigate the runtime and result quality of the proposed reduced-width ansatz design pattern. To facilitate a basic scaling analysis, we run all experiments with an iteratively increasing number of maximal layers up to four total.

5.1 Execution Times

We measure the time complexity in terms of the Average Execution Time (AET), which is comprised of the




Figure 5: Average execution time necessary to train each layer for the different circuit design schemes.

overall circuit execution and the time spent on training. For all circuit-designs and their different number of layers, the AET is measured in seconds over the 50 synthetically generated problem instances. The results are displayed in Fig. 5.

For the full-width design and the random-width design, we observe increasing AETs with increasing number of layers. This effect is stronger for the full-width design than for the random-width design. We explain those results by the two major tasks the simulator is performing: The circuit simulation and the parameter optimization. With increasing number of layers, the number of gates to simulate increases, which in turn leads to increasing circuit simulation times. However, the training time dominates the circuit execution time heavily as a closer inspection of the data shows (e.g., in layer one). Since the random-width design implements fewer gates overall and as the reduction of gates and parameters is distributed randomly across all layers, the AET curve has a flatter slope than the curve of the full-width design.

When analyzing the AET development of the reducing-width approach, we observe another pattern. In layer one, the AET coincides with that of the full-width design, as both designs implement the same circuit in the first layer. On the other hand, the reducing-width approach has a lower AET here, as its pruning already starts in layer one. From layer one to layer two, the AETs slightly increase for the reducing-width design, while for the remaining layers we observe significant decreases of its AETs. We can also explain this behavior by looking at the supposed development of the two main components of the AET. We assume, that the circuit simulation time increases throughout the layers at a decreasing pace, however, the absolute parameter optimization time decreases faster, as fewer gates are appended to the circuit and therefore fewer parameters need to be trained. While the effect of increasing circuit simulation time is still stronger in the step from the first layer to the second layer than the effect of decreasing optimisation time,




Figure 6: Average Approximation Ratio of the different circuit design schemes. Higher is better.

this ratio changes constantly, starting at the transition from layer two to three. The AET is almost halved from layer one to layer four. These results indicate good scalability properties of the reducing-width design in regards to training time.

5.2 Average Approximation Ratios

In order to assess the performance of different design schemes qualitatively, we investigate their Average Approximation Ratios (AARs) of the best possible graph partition. The optimal solution of every problem instance is calculated by brute force and the AAR describes the deviation of given solution from it. An AAR of 1 indicates perfect results, while an AAR of 0 is the worst possible solution.

Fig. 6 illustrates the evolution of the different AARs depending on the number of layers they implement. The plotted results show that the AARs generally decrease slightly with an increasing number of layers, but no clear systematic difference between the curves is apparent. In line with experience in QNN performance, further tests (not displayed here) without noise show the usually expected increase in performance for increasing circuit depth. Even though the noise level is quite small (merely a 1% error), related work on very similar problems has shown that including circuit noise in training can severely deteriorate the overall performance comparable to our findings (Borras et al., 2023; Wang et al., 2022; García-Martín et al., 2023; Stein et al., 2023). While the random-width design shows a slightly better ARR performance in layer four compared to the other two approaches, our small case study limits general conclusions on this advantage. Nevertheless, our data hints towards an advantage of the random pruning approach.

Combining these results with the runtimes, we can clearly see, that the pruning based circuit designs train substantially faster while achieving the same solution quality in the presence of noise.

6 CONCLUSION

In this article, we proposed the novel *reduced-width* ansatz design pattern, that uses gate pruning to improve the trainability of a QNN. Our experiments have shown that this ansatz design can significantly speed training time (up to 5 times) on noisy quantum simulators while achieving similar or even better solution quality. This inclines us to argue that current ansatz design schemes should be more open towards smaller/pruned circuits, as they show potential to exceed the performance of overparameterized circuits.

Future work must verify its applicability to a broader range of base ansätze as well as its scaling behavior for larger circuits and problem instances. To explore beyond our noisy test environment, noise free simulations should be carried out, while also exploring different parameter training approaches. In conclusion, we presented a potent, novel ansatz design pattern for quantum machine learning, that opens up a new path towards trainable yet expressive quantum neural networks.

ACKNOWLEDGEMENTS

This paper was partially funded by the German Federal Ministry of Education and Research through the funding program "quantum technologies - from basic research to market" (contract number: 13N16196). Furthermore, this paper was also partially funded by the German Federal Ministry for Economic Affairs and Climate Action through the funding program "Quantum Computing – Applications for the industry" (contract number: 01MQ22008A).

REFERENCES

- Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., and Woerner, S. (2021). The power of quantum neural networks. *Nat. Comput. Sci.*, 1(6):403–409.
- Borras, K., Chang, S. Y., Funcke, L., Grossi, M., Hartung, T., Jansen, K., Kruecker, D., Kühn, S., Rehm, F., Tüysüz, C., et al. (2023). Impact of quantum noise on the training of quantum generative adversarial networks. In *Journal of Physics: Conference Series*, volume 2438, page 012093. IOP Publishing.
- Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., and Coles, P. J. (2021a). Variational quantum algorithms. *Nat. Rev. Phys.*, 3(9):625–644.
- Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., et al. (2021b). Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644.

- Du, Y., Huang, T., You, S., Hsieh, M.-H., and Tao, D. (2022). Quantum circuit architecture search for variational quantum algorithms. *npj Quantum Inf.*, 8(1):62.
- Erdős, P. and Rényi, A. (1959). On random graphs i. *Publ. math. debrecen*, 6(290-297):18.
- Fösel, T., Niu, M. Y., Marquardt, F., and Li, L. (2021). Quantum circuit optimization with deep reinforcement learning. *arXiv preprint arXiv:2103.07585*.
- García-Martín, D., Larocca, M., and Cerezo, M. (2023). Effects of noise on the overparametrization of quantum neural networks. *arXiv preprint arXiv:2302.05059*.
- Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA. Association for Computing Machinery.
- Harrow, A. W., Hassidim, A., and Lloyd, S. (2009). Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103:150502.
- Hashim, A., Seritan, S., Proctor, T., Rudinger, K., Goss, N., Naik, R. K., Kreikebaum, J. M., Santiago, D. I., and Siddiqi, I. (2023). Benchmarking quantum logic operations relative to thresholds for fault tolerance. *npj Quantum Inf.*, 9(1):109.
- Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., and Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212.
- Hochreiter, S. (1998). The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Kobayashi, M., Nakaji, K., and Yamamoto, N. (2022). Overfitting in quantum machine learning and entangling dropout. *Quantum Machine Intelligence*, 4(2):30.
- Lanczos, C. (2012). *The Variational Principles of Mechanics*. Dover Books on Physics. Dover Publications.
- Liu, X., Angone, A., Shaydulin, R., Safro, I., Alexeev, Y., and Cincio, L. (2022). Layer vqe: A variational approach for combinatorial optimization on noisy quantum computers. *IEEE Transactions on Quantum Engineering*, 3:1–20.
- Lucas, A. (2014). Ising formulations of many np problems. *Frontiers in Physics*, 2.
- Martiel, S., Ayrat, T., and Allouche, C. (2021). Benchmarking quantum coprocessors in an application-centric, hardware-agnostic, and scalable way. *IEEE Transactions on Quantum Engineering*, 2:1–11.
- McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., and Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):4812.
- Mitarai, K., Negoro, M., Kitagawa, M., and Fujii, K. (2018). Quantum circuit learning. *Phys. Rev. A*, 98:032309.
- Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P. J., Aspuru-Guzik, A., and O'Brien, J. L. (2014). A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.*, 5(1):4213.
- Powell, M. J. (1994). *A direct search optimization method that models the objective and constraint functions by linear interpolation*. Springer.
- Scala, F., Ceschini, A., Panella, M., and Gerace, D. (2023). A general approach to dropout in quantum neural networks. *Advanced Quantum Technologies*, page 2300220.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117.
- Schuld, M., Bergholm, V., Gogolin, C., Izaac, J., and Killoran, N. (2019). Evaluating analytic gradients on quantum hardware. *Phys. Rev. A*, 99:032331.
- Schuld, M., Sweke, R., and Meyer, J. J. (2021). Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Phys. Rev. A*, 103:032430.
- Shor, P. (1994). Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134.
- Sim, S., Johnson, P. D., and Aspuru-Guzik, A. (2019a). Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070.
- Sim, S., Johnson, P. D., and Aspuru-Guzik, A. (2019b). Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070.
- Skolik, A., McClean, J. R., Mohseni, M., van der Smagt, P., and Leib, M. (2021). Layerwise learning for quantum neural networks. *Quantum Machine Intelligence*, 3:1–11.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958.
- Stein, J., Poppel, M., Adamczyk, P., Fabry, R., Wu, Z., Kölle, M., Nüßlein, J., Schuman, D., Altmann, P., Ehmer, T., et al. (2023). Quantum surrogate modeling for chemical and pharmaceutical development. *arXiv preprint arXiv:2306.05042*.
- Stougiannidis, P., Stein, J., Bucher, D., Zielinski, S., Linnhoff-Popien, C., and Feld, S. (2023). Approximate lookup-tables and arbitrary function rotations for facilitating nisq-implementations of the hhl and beyond. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 01, pages 151–160.
- Wang, H., Gu, J., Ding, Y., Li, Z., Chong, F. T., Pan, D. Z., and Han, S. (2022). Quantumnat: Quantum noise-aware training with noise injection, quantization and normalization. In *Proceedings of the 59th ACM/IEEE Design Automation Conference, DAC '22*, page 1–6, New York, NY, USA. Association for Computing Machinery.