KI - Künstliche Intelligenz
https://doi.org/10.1007/s13218-024-00862-9

PROJECT REPORTS

$Q(AI)^2$: Quantum Artificial Intelligence for the Automotive Industry

Tobias Stollenwerk[1] · Somtapa Bhattacharya[8] · Michele Cattelan[7] · Alessandro Ciani[1] · Gabriele Compostella[7] · David Headley[1,6] · Johannes Klepsch[5] · Matthias Klusch[3] · Markus Leder[6] · Antonio Macaluso[3] · Kristel Michielsen[2] · Dmytro Nabok[2] · Anestis Papanikolaou[8] · Alexander Rausch[4] · Marco Schumann[1] · Andrea Skolik[8] · Sheir Yarkoni[7] · Frank K. Wilhelm[1]

Received: 20 February 2024 / Accepted: 2 July 2024

1 Introduction

The goal of the project $Q(AI)^2$ was to acquire a broader basis of quantum computing enhanced AI and optimization algorithms for potential applications in the automotive industry. In particular, we strove to assess the potential for quantum advantage for specific use cases from the automotive industry. For this, we developed quantum algorithmic solutions tailored both to the available hardware as well as to the industrial problems at hand. In addition, we ensured industrial relevance of the investigated applications by the constitution of the consortium, involving all three large German car manufactures VW, Mercedes-Benz and BMW as well as Bosch as a central component supplier. Forschungszentrum Jülich and the German Center for Artificial Intelligence (DFKI) provided the necessary expertise in AI and quantum computing methods.

In this work we will report on the results of the above approach and provide an outlook to the future of quantum-accelerated AI applications in the automotive sector. We begin by reporting on our result on quantum supervised learning methods with a particular focus on quantum kernel methods in Sect. 2 applied to use cases from computational engineering and quality assurance. Next, we cover quantum reinforcement learning approaches applied to collision-free navigation of self-driving cars in Sect. 3, before we turn to the solution of planning and scheduling problems from the automotive industry, like flexible job-shop scheduling and ride-pooling with quantum optimization methods in Sect. 4. Lastly, we report on our findings concerning benchmarking use-cases and algorithms with a focus on the restrictions imposed by the inevitable noise of real quantum computer hardware on variational quantum algorithms in Sect. 5, before we conclude with a summary of our main lessons learned in the project.

2 Quantum Supervised Learning

In order to investigate the applicability of quantum-enhanced supervised machine learning methods to applications from the automotive industry, we took a three-pronged approach. First, we report on our work on tensor network quantum circuit, before we move to quantum kernel methods. Then, we cover classical-inspired quantum models.

2.1 Tensor Network Quantum Circuits

In this section, we consider tensor network quantum circuits. Designing circuits for quantum machine learning continues to present significant challenges. One potential solution that has recently gained attention is structuring the circuits based on tensor networks, a concept that has proven useful across multiple fields and particularly in classical machine learning.

✉ Tobias Stollenwerk
to.stollenwerk@fz-juelich.de

[1] Institute for Quantum Computing Analytics (PGI-12), Forschungszentrum Jülich, Jülich 52425, Germany

[2] Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich, Jülich 52425, Germany

[3] German Research Center for Artificial Intelligence (DFKI), Stuhlsatzenhausweg 3, Saarbrüecken 66123, Germany

[4] Robert Bosch GmbH, Corporate Research, Robert-Bosch-Campus 1, Renningen 71272, Germany

[5] BMW AG, Petuelring 130, München 80809, Germany

[6] Mercedes-Benz AG, Mercedesstraße 120, Stuttgart 70372, Germany

[7] Machine Learning Research Lab, Volkswagen Group, Munich 80805, Germany

[8] Volkswagen Data:Lab, Volkswagen Group, Munich 80805, Germany

Published online: 17 July 2024

A study we contributed to [1] provides an in-depth description of tensor-network quantum circuits and elucidates how to simulate their implementation. It incorporates a technique known as circuit cutting, which allows for the evaluation of circuits that contain a greater number of quantum bits (qubits) than those currently supported by existing quantum device technology. These methods were applied to the classification of welding defects.

To illustrate the computational necessities and potential applications, we performed simulations of various tensor-network quantum circuits using PennyLane [2], an openly available Python library that facilitates differential programming of quantum computers. Lastly, we provided demonstrations of these circuits' application to progressively complicated image processing tasks. The study can be seen as an exploration of a versatile method for circuit design applicable to tasks of industrial-scale machine learning. We observed that variational quantum circuits based on tensor network architectures can be used to detect welding defects (cf. Sect. 2.2) with good accuracy. However, additional comparisons to state-of-the-art classical methods are needed to assess whether there can be quantum advantage for this application.

2.2 Quantum Kernel Methods

The combination of quantum computing and kernel methods is a promising candidate for near-term quantum advantage and has provided a wide field for the design of quantum machine learning algorithms [3]. With the goal of determining whether kernel-based quantum machine learning methods can surpass classical ones on classification tasks based on actual data sets from the automotive industry in terms of classification results, we considered two challenging use cases.

The first use case was the quality estimation of *FEM meshes* as they typically occur in simulations of components based on the finite element method (FEM). For these, meshes of polygonal elements are created in an automated way whose quality is rated by domain experts as "good" or "bad" for potential manual rework. A supervised machine learning model is to learn these ratings based on geometric features of the elements in order to automatize the quality labelling [4].

The second use case we considered was similar to the first one but concerning *welding dots*. During the automated industrial welding processes, various physical parameters are recorded for each welding dot. For some of these dots, their quality is checked manually and labelled as "good" or "bad" in the training and test data. Note that due to the long history of refinements of the welding process, the recorded data sets are usually imbalanced and yield a low number of "bad" welding dots. A supervised machine learning model is to learn these labels based on said parameters as features to give hints on which newly welded dots to check predominantly. For these two cases, we considered three different quantum kernel methods.

The first method we investigated was Quantum Kernel Estimation (QKE) [5, 6] where a quantum circuit V is devised, encoding the feature vector x of a data point into a quantum state by having $V(x)$ act on the initial state with all qubits in state zero, $|0\rangle$. The kernel element $\kappa(x, y)$ is defined as the squared modulus of the inner product of two such states,

$$\kappa(x, y) = |\langle 0|V^{\dagger}(y)V(x)|0\rangle|^2 ,$$

which can be evaluated on a quantum computer. To determine the encoding circuits V and in turn the kernel κ, we varied several parameters in a combinatorial fashion: the number of repeated circuit encoding layers, the quantum gates out of which each layer is composed, and the scaling factor α for the feature vector x. As quantum gates, we used several data-encoding patterns inspired by [7].

As a second kernel method, we considered projected QKE (pQKE). This is similar to QKE, but after encoding the feature vector into a quantum state in the high-dimensional Hilbert state space, the state is projected into a lower-dimensional subspace to reduce its expressivity [8], as too much of it can hamper classification success [9].

Lastly, Quantum Kernel Training (QKT) [10] was applied as the third method. In QKT, based on the analogous classical method of kernel-target alignment [11], the defining circuit V is equipped with additional gates dependent on trainable parameters. We have used the best circuits gained from the QKE and added a layer of $R_y(\theta)$ gates right before it, where θ is only one trainable number for all gates (due to resource constraints). The kernel thus becomes

$$\kappa_{\theta}(x, y) = |\langle 0|R_y^{\dagger}(\theta)V^{\dagger}(y)V(x)R_y(\theta)|0\rangle|^2 .$$

For training, we have split the complete data set for each use case into natural subdivisions, e.g. the various FEM meshes. The total number of data points was about 6000 FEM elements and 5100 welding dots. Using a 80/20 split between training and testing data for each subdivision, we have performed a combinatorial grid search on the parameters mentioned above, and for each parameter combination a stratified 3-fold cross validation training with balanced accuracy as metric. The best-performing classifier was retrained on the whole training set of the respective subdivision and tested on its testing set.

In Table 1 we display the mean and standard deviation of balanced accuracy and the F1 score on the minority class calculated across the subdivisions. For comparison, the results of the classical RBF kernel are also reported where

Table 1 Results for three different quantum kernel methods and a classical reference method on several data sets of the use cases *FEM meshes* and *welding dots*

Use case	Balanced Accuracy (averaged)				F1 (minor class, averaged)			
	QKE	pQKE	QKT	RBF	QKE	pQKE	QKT	RBF
FEM	0.73 ± 0.13	0.73 ± 0.13	0.72 ± 0.14	0.75 ± 0.14	0.54 ± 0.25	0.53 ± 0.26	0.52 ± 0.28	0.56 ± 0.25
WD	0.66 ± 0.14	0.67 ± 0.14	0.65 ± 0.10	0.65 ± 0.13	0.57 ± 0.30	0.57 ± 0.32	0.55 ± 0.28	0.56 ± 0.30

the grid search was performed over the kernel parameter γ and an SVM regularization parameter C.

The results in Table 1 show that all three quantum methods range somewhat precisely on the same level as the classical reference SVM, irrespective of the encoding circuit and other parameters, although with a high standard deviation. For the FEM dataset, we have found with QKT a recall of 54% and precision of 63%. The best model for this problem in [4] was "extremely randomized trees" where, using a thresholding method for trading off, at a recall of 54% a precision of only ~50% was found, and at a precision of 63% a recall of only ~20%. Thus, our model might seem to outperform this. However, we could use only 6000 data points compared to other works [4] where 1.6 million have been used. Therefore, no meaningful comparison between the two methods can be made.

Altogether, quantum kernel methods were found to be applicable to industrial use cases and show results on par with off-the-shelf methods. However, no practical advantages were found for dataset sizes accessible to mere simulations of the quantum circuits. On theoretical grounds there is indication that the situation will be similar using larger data sets [12].

2.3 Classical-inspired Quantum Models

The aforementioned limitations of existing approaches indicate the necessity to propose novel methods to unlock the potential of quantum models in machine learning (ML). In this regard, an interesting direction involves constructing a quantum architecture capable of generating a quantum state equivalent to the output of classical supervised models and comparing the results with existing quantum models. Crucially, the objective of this research line is not to test quantum models in real-world use cases but rather to methodologically explore the properties of quantum models whose architecture is inspired by classical algorithms.

In this context, some of the authors proposed a universal and efficient framework [13] that reproduces the outputs of various classical supervised algorithms by leveraging the advantages of quantum computation. This framework can integrate multiple and diverse functions and has the potential to serve as the quantum counterpart for models falling under the category of aggregating multiple functions, such as ensemble algorithms and neural networks.

From a computational standpoint, the proposed framework facilitates the generation of an exponentially large number of different transformations of the input, with the increase in the depth of the corresponding quantum circuit occurring linearly. As specific instances, the quantum counterpart of the classical single-layer perceptron [14] and quantum ensemble [15] have been introduced. Nevertheless, employing quantum models to implement parametrized functions typically estimated by classical methods introduces the necessity of estimating non-linear functions through parametrized quantum circuits. This challenge has been partially addressed through the proposal of a quantum algorithm for spline functions [16].

These contributions have undergone experimental testing using standard data commonly employed in classical supervised learning problems. In some cases, they have outperformed kernel methods and quantum neural networks on these datasets. This underscores the need for further investigation and the proposal of novel quantum architectures to fully leverage Quantum Machine Learning (QML) methods. However, it is important to note that several technical modifications are necessary for these new approaches to be effectively adapted for industrial use cases, and these adjustments have been deferred to future work.

3 Quantum Reinforcement Learning

The application of reinforcement learning (RL) in automotive systems holds promise for enhancing autonomous driving capabilities, enabling vehicles to adapt and optimize their behavior based on real-time interactions with dynamic environments. This adaptive approach can contribute to improved decision-making and safety in complex driving scenarios. The recent adoption of quantum computing for RL has surfaced as an opportunity to surpass the limitations of classical methods by harnessing the potential advantages offered by quantum algorithms. In particular, the standard approach in quantum reinforcement learning (QRL) involves replacing a conventional deep neural network in a classical RL scenario with a parametrized quantum circuit to potentially enhance the learning capabilities of the agent.

A prevalent challenge in the current landscape is the reliance on benchmarks that involve fully observable and relatively simple environments that do not capture the complexity of real-world applications. In these scenarios, the quantum agent serves as an actor within the environment, a design choice that would necessitate a quantum computer during testing. As a result, the applicability of current QRL methods is very limited, highlighting the need for advancements that can handle more complex, partially observable, and diverse environments to broaden the practical scope of QRL. This limitation restricts the practical applicability of QRL methods and poses intriguing opportunities for future research.

Here, we present two algorithmic advances in quantum reinforcement learning achieved in $Q(AI)^2$. These advancements demonstrate the potential benefits of incorporating quantum computing into classical RL architectures, showcasing improved training performance for self-driving cars on real-world data using quantum simulation. Furthermore, a second investigation emphasizes how specific geometric properties of the data can be used to infer new ansätze with higher trainability and enhanced performance.

3.1 Quantum-supported Reinforcement Learning for Collision-Free Navigation of Self-Driving Cars

The challenge of collision-free navigation (CFN) for self-driving cars is an NP-hard problem that is usually tackled through deep reinforcement learning (DRL). In this regard, we developed a quantum-supported DRL algorithm for CFN in self-driving cars [17]. The method leverages quantum computation to enhance the training performance of DRL agents without requiring onboard quantum hardware. Based on the actor-critic approach, the method implements the critic using a hybrid quantum-classical algorithm suitable for near-term quantum devices. The performance was evaluated using the CARLA driving simulator, a benchmark for state-of-the-art DRL methods. Empirical assessments indicated that our method might outperform its classical counterpart not only in terms of training stability but also, in some instances, in terms of convergence rate when analyzing the *reward vs. episode* curve. This enhancement is achieved without adversely affecting the learned policy by the agent. Furthermore, indicated performance in terms of effective dimension, uncovering that including a quantum component might result in a model with greater descriptive power compared to classical baselines.

3.2 Equivariance Quantum Circuits

We designed a variational quantum algorithm that performs neural combinatorial optimization (NCO) by tuning the parameters of a quantum circuit with a classical optimization routine [18]. Quantum and classically, a major design decision in this type of algorithm is the architecture of the neural network (NN) or the parametrized quantum circuit (PQC). Classically, model architectures based on geometric deep learning have been widely successful. One special type of geometric model that has been used for combinatorial optimization problems on graphs is the graph NN. A graph neural network (GNN) takes as input a graph instance, and then aggregates node and edge information iteratively in each layer to learn a new embedding of the graph that can then be used to solve the problem at hand. On the quantum side, there have been some approaches to translate the GNN formalism to PQCs, however, most of those works have not taken into account geometric properties of the resulting models. Models in classical geometric deep learning have been shown to be successful because they respect certain symmetries that are present in the training data, e.g. invariance under node permutation in case of graphs, and therefore yield provably more data-efficient models. In our work, we developed a PQC for NCO on weighted graphs that is equivariant under node permutations and show that it widely outperforms more unstructured quantum models. To quantify the performance improvements of a symmetry-preserving model, we designed PQCs that gradually break the equivariance property, and show that the equivariant model outperforms all of them even on small problem instances. Additionally, we proved that even in the shallow regime, our ansatz is capable of producing the optimal solution for instances of arbitrary size, given the optimal setting of parameters can be found by the classical optimizer. Our work illustrates the fact that applications of QML can only be successful in the near-term if problem-tailored ansätze for PQCs are developed, and provides a successful ansatz for learning tasks on weighted graphs [19, 20].

4 Quantum Optimization

Quantum optimization is the application of quantum algorithms to solve optimization problems known to be hard to solve classically (typically belonging to the complexity class NP-hard). There are multiple paradigms to construct such quantum optimization algorithms. The best-known such algorithm is Grover's search, which exploits amplitude amplification as a subroutine to provide provable quantum speedups over classical algorithms. However, we focus on quantum annealing (QA) [21] and the quantum approximate optimization algorithm (QAOA) [22], as these methods are amenable to near-term devices.

In the following we report on our efforts to solve two applications from the automotive industry, bin packing and the ride-pooling problem, with the above methods. Afterwards, we focus on quantum assisted solutions to coalition formation in multi-agent systems and flexible job-shop scheduling, which are both applicable to a wider variety of optimization problems from the automotive industry. Finally, we report on an advanced method for handling hard constraints, as they typically occur in combinatorial optimization problems from industry.

4.1 Bin Packing

Bin packing is an NP-hard problem that plays a crucial role in optimizing resource allocation and storage, making it particularly important to automotive manufacturing where efficient space utilization directly impacts operational efficiency. However, it poses significant challenges in finding efficient solutions using state-of-the-art classical algorithms. In this respect, we developed a formulation as quadratic unconstrained binary optimization (QUBO) that utilizes the augmented Lagrangian method to seamlessly integrate bin packing constraints into the objective function, concurrently enabling an analytical estimation of heuristic yet empirically robust penalty multipliers [23]. This approach establishes a more versatile and generalizable model, eliminating the necessity for empirically calculating instance-dependent Lagrangian coefficients-a common requirement in alternative QUBO formulations for analogous problems. The experimental findings using a real D-Wave device not only validate the correctness of the proposed formulation but also showcase the potential of our approach in effectively addressing the bin-packing problem, particularly with the evolution of more reliable quantum technology.

4.2 Ride-Pooling problem

The ride-pooling problem (RPP), is a combinatorial optimization problem centered around the efficient coordination and scheduling of shared transportation services [24]. In this problem, a fleet of vehicles is tasked with serving a set of users who request rides between specified origins and destinations, each with individual time constraints. The primary objective is to optimize the allocation of vehicles to passengers, considering constraints such as vehicle capacity, time windows, and minimizing total travel distance or time. The goal is to enhance the overall efficiency of transportation services by facilitating shared rides, thereby reducing congestion, fuel consumption, and overall operational costs. The RPP finds application in urban mobility, public transportation, and emerging ride-sharing services, where the aim is to enhance the sustainability and effectiveness of transportation systems. The investigation of RPP has gained increased interest in recent years through the emergence of car-sharing services such as Uber and Lyft, and is therefore of practical interest to both model and solve accurately in practice. We focused on solving the RPP using distance minimization of vehicles in the fleet to allow a straightforward motivation (and comparison) with alternative QUBO models. The basic methods developed in this paper can extend to almost any of these objective functions, and therefore are general enough to be used for a wide variety of applications. Specifically, we showed how to construct QUBOs that accurately represent the RPP problem through the inclusion of problem-specific constraints, compare this to canonical representations, and estimate the resources required to solve such RPP QUBOs using quantum processors and quantum algorithms [25]. We demonstrated how the methods we developed could lead to more resource efficient problem representations for quantum optimiziation algorithms.

4.3 Coalition Formation in Multi-Agent Systems

Coalition formation in multi-agent systems [26] has diverse applications in the automotive industry. It enables autonomous agents, representing vehicles or components, to form alliances and collaboratively address challenges. Examples include optimizing traffic flow, managing fleets efficiently, coordinating energy-efficient driving, ensuring cooperative collision avoidance, streamlining supply chains in manufacturing, facilitating autonomous vehicle coordination, and enhancing shared mobility services. In this area, three pioneering quantum-supported algorithms have been developed to address the intricate coalition structure generation (CSG) problem. Each algorithm provides distinctive contributions aimed at harnessing the power of quantum computing.

We developed a method for addressing the CSG problem by reformulating it as a QUBO [27]. This approach facilitates the determination of optimal coalition structures through established quantum techniques, including QAOA and Quantum Annealing, applied to the most general formulation of coalition games and demonstrating potential advantages over classical baselines in terms of runtime. We performed small-scale experiments using IBM Qiskit and D-Wave machines, confirming its performance across various scenarios. Nevertheless, the algorithm displays elevated computational demands, as the required number of qubits scales exponentially with the number of agents, rendering it non-friendly to near-term quantum computers.

To address these limitations, two alternative quantum-supported solutions have been proposed. These alternatives concentrate on a specific category of games known as induced subgraph games, where coalition games are induced by connected, undirected weighted graphs [28]. First, [29] addresses the CSG by employing a quantum annealing device to iteratively split coalitions into two nonempty subsets, maximizing the coalition value. This quantum-supported approach exhibits good performance compared to state-of-the-art classical solvers, showcasing a runtime quadratic in the number of agents and a worst-case approximation ratio of 93%. Second, [30] adopts the same hybrid quantum-classical strategy of the former method, by reformulating the partition problem as a QUBO and solving it with QAOA. Still, this method not only demonstrates superior performance in terms of runtime and approximation ratio but also requires fewer qubits with respect to our original approach, enabling experiments on medium-sized problems. Collectively, these quantum approaches represent significant strides toward addressing the CSG problem, offering novel insights into the potential of quantum computing in optimizing rational agent coalition games. The experimental validations underscore the efficiency and effectiveness of these algorithms, laying the groundwork for further exploration and application in complex problem-solving domains.

Finally, the same algorithmic approach has recently been explored for quantum-supported unsupervised image segmentation [31]. In this instance, the utilization of quantum annealing for performing graph-cutting on a graph generated from images has demonstrated notable outcomes, surpassing current state-of-the-art optimizers in terms of runtime and worst-case approximation, achieving up to 95%. These results hold even when considering QUBO problems with up to 2000 logical variables.

4.4 Flexible Job Shop Scheduling

The challenges of flexible job shop scheduling (FJSS) resonate significantly in manufacturing and automotive industries, where optimized resource allocation and efficient task scheduling are crucial for enhancing production workflows. The FJSS problem is renowned for its inherent complexity, characterized by the need to optimize the allocation of resources and tasks across multiple machines with varying processing times. Conventional methods often face challenges in providing efficient solutions due to the problem's NP-hard nature. In the quest to address this complexity, quantum algorithms have emerged as potential candidates [32].

Notably, all the proposed quantum formulations introduce the concept of a *timeline* to facilitate the reformulation as a QUBO. However, this formulation imposes a substantial number of logical variables. To refine this approach, we explored a time-independent formulation that eliminates the concept of a timeline and facilitates the systematic derivation of fewer logical variables compared to other prevailing (quantum) methods. With this novel formulation, we were able to test small-problem instances to confirm the correctness of the proposed approach in providing near-optimal solutions. Despite these advancements, addressing real-world industrial instances remains impractical with current quantum technology. Nevertheless, considering that state-of-the-art solutions for the FJSS problem predominantly rely on genetic algorithms, we posit that this avenue holds promise for advancing solutions to the challenging task of FJSS.

4.5 Parallel Implementation of Variational Quantum Algorithms

We developed a method for decomposing quantum circuits of variational quantum algorithms (VQA) utilized to implement combinatorial optimization problems into smaller circuits amenable to parallel training [33]. Our method is based on the observation that the space of solutions of a combinatorial optimization problem can often be interpreted as Cartesian products of vector spaces. To illustrate, in the context of an optimization problem featuring two constraints, the set of feasible solutions is, at most, the product of the sets of feasible solutions for each constraint independently.

To apply our method we reformulated the optimization problem as independent slices, with each slice representing a sub-problem obtained by removing a constraint. For constrained optimization problems the construction of such slices can be done starting from the definition of the variables of the problem and, therefore, can be done a priori without having any knowledge about the quantum circuit. We collected the results from each smaller quantum circuit and we updated the parameters of the VQA using the complete problem definition, thereby enforcing the inclusion of the 'missing' constraint. The solutions used to train the VQA are retrieved by multiplying the solutions from individual slices.

It is worth noticing that for various constrained optimization problems, the independent slices are essentially identical copies of each other. Consequently, it becomes unnecessary to simulate all slices; instead, simulating a single slice and taking the product of solutions with itself suffices as candidate solutions for the complete problem. For instance, we considered QAOA, its parallelized version as described above and the implementation of the former by using a single slice. We noticed that the single-slice implementation exhibits a polynomial reduction in resource requirements compared to both standard QAOA and its direct parallelization. Our experiments revealed that employing the parameters trained by implementing a single slice in regular QAOA circuits yields results comparable to training the complete circuit with a significantly larger number of qubits. Thus, our findings suggest the existence of redundancy in encoding the full problem Hamiltonian in QAOA, with not all components essential for computing relevant solutions to the original optimization problem.

In conclusion, we can see that our parallel implementation can match the performance of QAOA by using factorially fewer samples to train the hyperparameters. However, the effort required by the classical optimizer is increased and rules to select the slices must be chosen. Furthermore, more extensive experiments on real hardware or simulated noise must be conducted to validate the practicability of this parallelization method.

5 Implementation and Benchmarking

Due to the inherent error prone nature of near-term quantum computing devices and the corresponding algorithms, hardware restrictions play a crucial, potentially inhibitive role for the performance of the outlined quantum algorithms for applications from the automotive industry. In addition, the heuristic nature of these approaches demand for a broad data base of representative test problems. Therefore, careful benchmarking and the consideration of noise etc., are necessary to truly assess the capabilities of quantum computers for applications in the automotive industry. In this section, we report on our efforts towards these topics in Q(AI)^2.

5.1 Benchmarking Framework

We contributed to the quantum application benchmarking framework QUARK [34, 35], which is covered in details in another article of this issue.

5.2 Optimal Parameters for QAOA

When using QAOA to solve combinatorial optimization problems, finding parameters that maximize the expected value of solutions can present a barrier to its effective use. It is, therefore, of value to explore mathematical methods for determining optimal parameters by exploiting coarse-grained properties of a problem or class of problems one wants to solve, without incurring the cost of either an iterative loop-optimization procedure on a noisy quantum processor or the inefficient classical simulation thereof.

While prior work in this direction [36, 37] has focused on computing the expected solution quality for specific problem classes, we take the approach of modelling problems as random variables with a given distribution of costs [38]. Through this formulation, and in combination with the use of a highly symmetric driver as used in Grover's algorithm, it becomes possible to write a closed form expression for the expected solution quality, without a computationally prohibitive dependence on the size of the problem solved. Work in this direction could help the applicability of QAOA to industry problems by the provision of variational parameters to use, or start some optimization with, based on statistical properties of the problem in question.

5.3 Noise-induced Barren Plateaus

The performance of variational quantum algorithms depends on the ability of the classical optimizer to find suitable circuit parameters. Under certain conditions the gradient of the cost function with respect to the circuit parameters vanishes exponentially. This phenomenon is called a barren plateau and reduces the efficiency of the optimization process. Therefore it is crucial to understand the circumstances that lead to barren plateaus. For instance, it has been shown that highly expressive circuits induce barren plateaus [39], as does local Pauli noise [40]. Whenever noise is the reason for a barren plateau, it is referred to as a *noise-induced barren plateau* (NIBP).

In Q(AI)^2, we studied which other noise models give rise to NIBP [41]. To this end, we used layered noise models. In each layer ℓ we applied a parameterized unitary $U(\boldsymbol{\theta}_{\boldsymbol{\ell}})$ depending on K parameters gathered in the vector $\boldsymbol{\theta}_{\boldsymbol{\ell}} = (\theta_{\ell 1}, \ldots, \theta_{\ell K})$ followed by a noise channel $\mathcal{N}$. We analytically proved the existence of NIBP when (i) a parameter shift rule holds for the gradient and (ii) the unitaries are given by random unitary 2-designs. Assuming $\mathcal{N}$ to be a completely positive trace preserving map [42], we show that the gradient $\frac{\partial C}{\partial \theta_{\ell k}}$ vanishes exponentially in $L - \ell$. Considering the specific case of unital or weak Markovian noise in (ii) we find that $\frac{\partial C}{\partial \theta_{\ell k}}$ vanishes exponentially in L. As a more practical case, we consider QAOA circuits for MaxCut problems on d-regular graphs under weak amplitude damping noise. We find that the behavior of the QAOA circuits is well described by the toy model circuit, used in (ii). Thus, we also find indications of NIBP in the considered QAOA circuits. In summary, we found that NIBP are a significant challenge for the scalability of variational quantum algorithms as they are employed in Q(AI)^2, and further advances in algorithmic design seem necessary to push the field forward.

6 Summary

In this project, we have investigated algorithms such as quantum enhanced machine learning and optimization algorithms for potential applications in the automotive industry. Our goal was to gain a better understanding of the value that quantum computing may bring to the automotive industry in the future. We found that although there is potential for quantum advantage with regard to applications from the automotive industry, it is very challenging to assess future performance due to the limited capabilities of near-term quantum computing devices. However, the results obtained during this project are, in general, encouraging. Additionally, we found that expertise gained by implementation of concrete use cases into quantum algorithms can help any company to prepare for a future breakthrough in quantum computing technology. An overview of the scientific articles produced during the project can be found on our webpage [43]. The constitution of the project, in particular the open collaboration of the industrial research groups in a pre-competitive environment, as well as the involvement of focused research groups with expertise in quantum and AI methods, proved to be very fruitful.

Acknowledgements The authors acknowledge the support of the German Federal Ministry of Education and Research (BMBF) in the funding program "quantum technologies - from basic research to market" for the project $Q(AI)^2$ – Quantum Artificial Intelligence for the Automotive Industry.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

1. Guala D, Zhang S, Cruz E, Riofrío CA, Klepsch J, Arrazola JM (2023) Practical overview of image classification with tensor-network quantum circuits. Sci Rep 13(1):4427
2. Bergholm V, Izaac J, Schuld M, Gogolin C, Ahmed S, Ajith V, Alam MS, Alonso-Linaje G, AkashNarayanan B, Asadi A, Arrazola JM, Azad U, Banning S, Blank C, Bromley TR, Cordier BA, Ceroni J, Delgado A, Matteo OD, Dusko A, Garg T, Guala D, Hayes A, Hill R, Ijaz A, Isacsson T, Ittah D, Jahangiri S, Jain P, Jiang E, Khandelwal A, Kottmann K, Lang RA, Lee C, Loke T, Lowe A, McKiernan K, Meyer JJ, Montañez-Barrera JA, Moyard R, Niu Z, O'Riordan LJ, Oud S, Panigrahi A, Park CY, Polatajko D, Quesada N, Roberts C, Sá N, Schoch I, Shi B, Shu S, Sim S, Singh A, Strandberg I, Soni J, Száva A, Thabet S, Vargas-Hernández RA, Vincent T, Vitucci N, Weber M, Wierichs D, Wiersema R, Willmann M, Wong V, Zhang S, Killoran N (2018) Pennylane: automatic differentiation of hybrid quantum-classical computations
3. Schuld M, Petruccione F (2021) Machine learning with quantum computers. In: Quantum Science and Technology. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-030-83098-4
4. Sprave J, Drescher C (2021) arXiv:2107.10507 [cs, math]
5. Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM (2019) Supervised learning with quantum-enhanced feature spaces. Nature 567(7747):209. https://doi.org/10.1038/s41586-019-0980-2
6. Schuld M, Killoran N (2019) Quantum machine learning in feature hilbert spaces. Phys Rev Lett 122(4):040504. https://doi.org/10.1103/PhysRevLett.122.040504
7. Sim S, Johnson PD, Aspuru-Guzik A (2019) Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. Adv Quant Technol 2(12):1900070. https://doi.org/10.1002/qute.201900070. https://onlinelibrary.wiley.com/doi/abs/10.1002/qute.201900070
8. Huang HY, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, McClean JR (2021) Power of data in quantum machine learning. Nat Commun 12(1):2631. https://doi.org/10.1038/s41467-021-22539-9
9. Thanasilp S, Wang S, Cerezo M, Holmes Z (2022) Exponential concentration and untrainability in quantum kernel methods . https://doi.org/10.48550/arXiv.2208.11060. arXiv:2208.11060
10. Glick JR, Gujarati TP, Corcoles AD, Kim Y, Kandala A, Gambetta JM, Temme K (2021) arXiv:2105.03406
11. Cristianini N, Kandola J, Elisseeff A, Shawe-Taylor J (2006) Innovations in Machine Learning pp. 205–256 . https://doi.org/10.1007/3-540-33486-6_8. https://link.springer.com/chapter/10.1007/3-540-33486-6_8
12. J. Kübler, S. Buchholz, B. Schölkopf (2021) Advances in Neural Information Processing Systems, vol. 34. Curran Associates, Inc., pp. 12661–12673. https://proceedings.neurips.cc/paper/2021/hash/69adc1e107f7f7d035d7baf04342e1ca-Abstract.html
13. Macaluso A, Klusch M, Lodi S, Sartori C (2023) MAQA: a quantum framework for supervised learning. Quant Inform Process 22(3):159
14. Macaluso A, Orazi F, Klusch M, Lodi S, Sartori C (2022) International conference on machine learning, optimization, and data science, Springer, pp. 341–356
15. Macaluso A, Clissa L, Lodi S, Sartori C (2024) IET Quant Commun:1–16 .https://doi.org/10.1049/qtc2.12087
16. Inajetovic MA, Orazi F, Macaluso A, Lodi S, Sartori C (2023) Computational Science—ICCS 2023. Springer, Cham, pp. 177–192
17. Sinha A, Macaluso A, Klusch M (2023) arXiv preprint arXiv:2311.12875
18. Cappart Q, Chételat D, Khalil EB, Lodi A, Morris C, Veličković P (2023) Combinatorial optimization and reasoning with graph neural networks. J Mach Learn Res 24(130):1
19. Skolik A, Cattelan M, Yarkoni S, Bäck T, Dunjko V (2023) NPJ Quant Inform 9(1):47
20. Skolik A, Mangini S, Bäck T, Macchiavello C, Dunjko V (2023) Quantum multi-agent reinforcement learning for aerial ad-hoc networks. EPJ Quant Technol 10(1):1
21. Farhi E, Goldstone J, Gutmann S, Sipser M (2000) arXiv:quant-ph/0001106
22. Farhi E, Goldstone J, Gutmann S (2014) arXiv preprint arXiv:1411.4028
23. Cellini L, Macaluso A, Lombardi M (2023) arXiv preprint arXiv:2309.12678
24. Ke J, Yang H, Zheng Z (2020) On ride-pooling and traffic congestion. Transp Res Part B Methodol 142:213
25. Cattelan M, Yarkoni S (2022) arXiv preprint arXiv:2212.04894

26. Rahwan T, Michalak TP, Wooldridge M, Jennings NR (2015) Coalition structure generation: a survey. Artif Intell 229:139. https://doi.org/10.1016/j.artint.2015.08.004
27. Venkatesh SM, Macaluso A, Klusch M (2022) Proceedings of the 19th ACM International Conference on Computing Frontiers (Association for Computing Machinery, New York, NY, USA, 2022), CF '22, p. 189–192. https://doi.org/10.1145/3528416.3530235
28. Voice T, Polukarov M, Jennings NR (2012) Coalition structure generation over graphs. J Artif Intell Res 45:165
29. Venkatesh SM, Macaluso A, Klusch M (2023) Computational Science—ICCS 2023. Springer, Cham, pp. 138–152
30. Venkatesh SM, Macaluso A, Klusch A (2023) Proceedings of the 20th ACM International Conference on Computing Frontiers (Association for Computing Machinery, New York, NY, USA, 2023), CF '23, p. 197–200. https://doi.org/10.1145/3587135.3592192
31. Venkatesh SM, Macaluso A, Nuske M, Klusch M, Dengel A (2023) arXiv preprint arXiv:2311.12912
32. Denkena B, Schinkel F, Pirnay J, Wilmsmeier S (2021) Quantum algorithms for process parallel flexible job shop scheduling. CIRP J Manuf Sci Technol 33:100. https://doi.org/10.1016/j.cirpj.2021.03.006
33. Cattelan M, Yarkoni S (2023) arXiv preprint arXiv:2304.03037
34. Finzgar JR, Ross P, Holscher L, Klepsch J, Luckow A (2022) 2022 IEEE International Conference on Quantum Computing and Engineering (QCE) (IEEE, 2022). https://doi.org/10.1109/qce53715.2022.00042
35. Kiwit FJ, Marso M, Ross P, RiofrÃo CA, Klepsch J, Luckow A (2023)Application-oriented benchmarking of quantum generative learning using quark
36. Farhi E, Goldstone J, Gutmann S, Zhou L (2022) The quantum approximate optimization algorithm and the Sherrington–Kirkpatrick model at infinite size. Quantum 6:759. https://doi.org/10.22331/q-2022-07-07-759
37. Claes J, Dam Wv (2021) Instance independence of single layer quantum approximate optimization algorithm on mixed-spin models at infinite size. Quantum 5:542. https://doi.org/10.22331/q-2021-09-15-542
38. Headley D, Wilhelm FK (2023) Revealing cell vulnerability in Alzheimer's disease by single-cell transcriptomics. Phys Rev A 107:012412. https://doi.org/10.1103/PhysRevA.107.012412
39. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H (2018) Barren plateaus in quantum neural network training landscapes. Nat Commun 9(1):4812. https://doi.org/10.1038/s41467-018-07090-4
40. Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L, Coles PJ (2021) Noise-induced barren plateaus in variational quantum algorithms. Nat Commun 12(1):6961. https://doi.org/10.1038/s41467-021-27045-6
41. Schumann M, Wilhelm FK, Ciani A (2023) arXiv preprint arXiv:2310.08405
42. Nielsen MA, Chuang IL (2011) Quantum computation and quantum information: 10th anniversary edition, 10th edn. Cambridge University Press, Cambridge
43. Q(AI)2 homepage. https://qai2.de. To be online on March 31