

# Quantum Artificial Intelligence: A Brief Survey

Matthias Klusch · Jörg Lässig · Daniel Müssig ·  
Antonio Macaluso · Frank K. Wilhelm

**Abstract** Quantum Artificial Intelligence (QAI) is the intersection of quantum computing and AI, a technological synergy with expected significant benefits for both. In this paper, we provide a brief overview of what has been achieved in QAI so far and point to some open questions for future research. In particular, we summarize some major key findings on the feasibility and the potential of using quantum computing for solving computationally hard problems in various subfields of AI, and vice versa, the leveraging of AI methods for building and operating quantum computing devices.

**Keywords** Quantum AI · Quantum Computing · AI

## Mathematics Subject Classification (2020)

68Q12: Quantum algorithms · 68T01: General AI

---

Matthias Klusch  
German Research Center for Artificial Intelligence GmbH  
(DFKI), Saarbruecken, Germany  
E-mail: matthias.klusch@dfki.de

Jörg Lässig  
University of Applied Sciences Zittau/Görlitz, Görlitz, Germany  
E-mail: jlaessig@hszg.de

Daniel Müssig  
Fraunhofer Institute of Optronics, System Technologies and  
Image Exploitation - Advanced System Technology (IOSB-  
AST), Görlitz, Germany  
E-mail: daniel.muessig@iosb-ast.fraunhofer.de

Antonio Macaluso  
German Research Center for Artificial Intelligence GmbH  
(DFKI), Saarbruecken, Germany  
E-mail: antonio.macaluso@dfki.de


Frank K. Wilhelm  
Forschungszentrum Juelich GmbH (FZJ), Jülich, Germany  
E-mail: f.wilhelm-mauch@fz-juelich.de

## 1 Introduction

It is known that quantum computing can simulate and even go beyond classical computing in terms of computational speedup in theory [3, 130, 153] for quite some time. But initial versions of real quantum computing hardware and frameworks for quantum programming became available only in about the past decade. Even the Quantum Internet of networked quantum computers and with secure quantum communication channels, long time considered as mere science fiction, is on its way with early stage prototypes available [26, 59, 33]. On the other hand, artificial intelligence (AI) [156] is commonly considered as one of the most disruptive key technologies of our time for industry and business, our private and social life, notwithstanding the challenges of its future, trustworthy and controlled use for the benefit of the people affected by it.

*Quantum Artificial Intelligence* (QAI, Quantum AI) is the intersection of both technologies (cf. Fig. 1) and concerned with the investigation of the feasibility and the potential of leveraging quantum computing for AI, and vice versa, AI for quantum computing [152]. While quantum machine learning [17, 16] is currently the most popular application [162, 47, 53, 117], quantum AI goes much beyond covering more subfields of AI [156, 73], such as quantum reasoning [36, 34, 21], quantum automated planning and scheduling (QPS), quantum natural language processing (QNLP), quantum computer vision (QCV), and quantum agents and multi-agent systems (QMAS). Notably, each of these QAI subfields covers research and development in both directions; for example, QML refers to both the use of quantum computing for machine learning and vice versa. In this respect, it would be premature to follow the current hype cycle around QML.

Though QAI still is a nascent and inherently interdis-



**Fig. 1** Quantum AI (QAI) as intersection of quantum computing and AI with subfields in relation to AI each covering both directions.

disciplinary research field, remarkable progress has been already made in the recent past. It has been shown, in part also experimentally on real quantum computers of the current NISQ (Noisy Intermediate-Scale Quantum) era, that tailored QAI algorithms could indeed make a positive difference in solving certain computationally hard problems such as for combinatorial optimization in AI applications compared to classical solutions. So far, the potential quantum utility of direct quantum or hybrid quantum-classical methods of QAI is under investigation for applications in diverse domains such as manufacturing, automated driving, transport and logistics, energy management, healthcare, finance, aerospace, climate and earth sciences, and pharmaceutical and chemical industries. In this regard, quantum AI use cases are concerned with, for example, portfolio optimization in finance [150], traffic management [193], capacitive vehicle routing [62, 82], safe navigation of self-driving vehicles [163], satellite network constellation [178] and mission planning [147], energy network management [22, 124], job-shop scheduling in manufacturing [158], weather simulation and forecasting [83, 68, 132, 170] and simulation of materials and chemicals for drug discovery [142, 91, 111] and material design [37, 10]. On the other hand, there is active research on the use of QAI algorithms that leverage AI methods, in particular from machine learning, to address challenges of the building and operation of quantum computing devices. These research activities in both directions of QAI are in part also fueled and influenced by the made progress and still ongoing race in building ever more powerful quantum computing devices and network, and vice versa. Currently, the global market value of QAI applications


on future quantum computing devices in general is estimated to be about eighteen billion USD by 2030 [148], and three to five billion USD for the automotive industry by 2035 in particular [169, 25]. This emphasizes that quantum AI has graduated from being a mere academic niche to a topic with a potential future beyond the current hype. However, it remains unclear when QAI methods and their applications, including those mentioned above, can be used and commercialized at large in practice, as this would require way more quantum computational resources and fault tolerance than current quantum computers have but future computing devices could provide [93].

Remarkably, as in AI, there are also discussions related to QAI on ethical issues of quantum computing such as those concerned with non-transparency of information processing in a “quantum-box”, or clashes of quantum privacy with security demands [139], as well as on the role of quantum computing in neuroscience (quantum neuroscience). The latter includes speculative multi-scale simulations of the human brain [168] and explanation of some consciousness-related brain functions [90], though not of how the human brain produces thoughts [108], as well as a better representation of and inference means for certain psychological models in cognitive science (quantum cognition) [140]. One should keep in mind, that these concerns apply in similar form to classical stochastic algorithms as well.

In this paper, we provide a first overview of selected methods, use cases and insights from research in the interdisciplinary field of QAI for both its directions, without any claim to completeness. In Section 2, we informally recall the basics of quantum computing from the computer science perspective only in very brief; readers who are roughly familiar with them can easily skip this section. Section 3 then summarizes selected key findings of research on quantum computing for AI in several subfields of QAI. The same is done for research on AI for quantum computing in Section 4 before we conclude in Section 5.

## 2 Quantum Computing in a Nutshell

Quantum computing [130, 153] harnesses the principles of quantum mechanics [149] to process information and perform computations, potentially surpassing the capabilities of classical computers. There are two primary models of quantum computing. The first is gate-based quantum computing, which functions analogously to classical computing by using quantum gates to manipulate quantum information. This model facilitates the design of complex quantum circuits, which are the



**Fig. 2** Representation of a quantum state  $|\psi\rangle$  on the Bloch sphere. The state is described by the angles  $\theta$  and  $\varphi$ , where  $\theta$  defines the polar angle from the  $z$ -axis and  $\varphi$  defines the azimuthal angle in the  $xy$ -plane. The Bloch sphere provides a geometric representation of the pure states of a qubit, with  $|0\rangle$  and  $|1\rangle$  corresponding to the poles of the sphere.

quantum counterparts to Boolean circuits in classical computation.

In contrast, adiabatic quantum computing [7] is based on the adiabatic theorem and, in principle, equivalent to quantum gate-based computing [4]. In this model, the quantum system, which represents the search space of the given computational problem, gradually evolves from an initial, simple quantum state to a final state that encodes the problem's solution. The choice between these models depends on the nature of the problem being addressed and the availability in terms of hardware.

## 2.1 Gate-Based Quantum Computing

The basic unit of information in classical information processing is a single bit that can exist in one of two states, represented as the integer numbers 0 or 1. Consequently, a sequence of  $n$  bits can represent  $2^n$  unique values, with the bit register being in only one of these  $2^n$  possible states. A *quantum bit* (qubit) is the quantum analog of a bit and can assume the two basic states  $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$  and  $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ . Qubits adhere to the principles of quantum mechanics and can be implemented using various physical systems, such as the spin states of subatomic particles, ion traps, neutral atoms, or superconducting circuits.

One of the intriguing properties of qubits is their ability to exist not just in the state  $|0\rangle$  or  $|1\rangle$  but in a *superposition* of both. An arbitrary *single-qubit state* can be expressed as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \quad (1)$$

where the coefficients  $\alpha$  and  $\beta$  are complex numbers that satisfy the normalization condition  $|\alpha|^2 + |\beta|^2 = 1$ . A visual representation of a qubit is shown in Figure 2. When building *multi-qubit systems*, an  $n$ -qubit system provides access to a  $2^n$ -dimensional Hilbert space, where an arbitrary pure quantum state is defined as

$$|\psi\rangle = c_0|0\cdots 0\rangle + c_1|0\cdots 1\rangle + \cdots + c_{2^n-1}|1\cdots 1\rangle, \quad (2)$$

with  $c_i \in \mathbb{C}$  and  $\sum_{i=0}^{2^n-1} |c_i|^2 = 1$ . The multi-qubit basis states, e.g.,  $|0\cdots 1\rangle = |0\rangle \otimes |0\rangle \otimes \cdots \otimes |0\rangle \otimes |1\rangle$ , are tensor products of the individual qubits. The state  $|\psi\rangle \rightarrow [c_0, c_1, \dots, c_{N-1}]^t$  possesses  $N = 2^n$  complex amplitudes, whose absolute squared values must sum to one. Thus, due to the principle of superposition, an  $n$ -qubit system is capable of encoding information that scales as  $2^n$ , while classical systems are limited to  $n$ . A multi-qubit state is *entangled* if the states of its component qubits cannot be described independently of each other as a tensor product. Quantum entanglement allows for non-local correlations such that operating on one component qubit changes the states of others instantaneously and independent of their separate locations; a prominent example are the entangled 2-qubit Bell states.


In order for computation to be possible, there must be a way to manipulate quantum states. This is achieved by operators that describe the evolution of a closed quantum system with unitarity. In order to preserve the normalization, these have to be unitary, and therefore reversible, that is, for a *quantum operator*  $U$  must hold that  $U^\dagger U = I$ . In fact, any operation on qubits can be described as a matrix operator. Notice that this is fundamentally different from classical computing, where operations are not required to be unitary and not even reversible.

Another crucial difference between classical and quantum computation is how information can be accessed after processing. In classical computing, the state of each bit is well-defined and can be directly observed at any time, revealing whether it is in a state of 0 or 1. This observation does not alter the state of the bits, allowing us to access the exact information stored in the system without any disturbance. However, as mentioned above, in quantum systems the state of a qubit is generally a superposition of multiple states, meaning it can exist in a combination of both 0 and 1 simultaneously. To extract information from a quantum system, one must measure an observable physical quantity - which from now on we assume to be the binary number associated with the computational basis states. This *measurement process* forces the quantum system to *collapse* from its superposition into one of the possible definite states. Crucially, this collapse is proba-

bilistic, meaning that the outcome of the measurement cannot be precisely predicted but rather is governed by the amplitudes of the quantum state. Furthermore, the measurement fundamentally alters the state of the system, making it impossible to retrieve the original superposition state after the measurement is performed. This is a key distinction from classical systems, where observation is non-intrusive and reversible.

For the quantum state given by Eq. (1), the measurement will yield one of the basis states  $|0\rangle$  or  $|1\rangle$  with classical output 0 or 1, respectively. The probability of observing  $|0\rangle$  is  $|\alpha|^2$ , and the probability of observing  $|1\rangle$  is  $|\beta|^2$ , with the total probability summing to 1 ( $|\alpha|^2 + |\beta|^2 = 1$ ). After measurement, the qubit will be found in the state corresponding to the measurement outcome. For the state described in Eq. (2), the measurement outcome will be one of the  $2^n$  possible basis states. The probability of observing a particular multi-qubit state  $|\psi\rangle = |b_1 b_2 \dots b_n\rangle$  is given by  $|c_{\text{index}}|^2$ , where index represents the binary number corresponding to the combination  $b_1 b_2 \dots b_n$ , and the total probability is 1 ( $\sum_{i=0}^{2^n-1} |c_i|^2 = 1$ ). After the measurement, the system collapses to the state observed, with the measurement outcome determining the final state of the system. Quantum state measurement is the only non-unitary (irreversible) operation in an ideal quantum computer, and according to the no-cloning theorem of quantum computing, in contrast to classical computing, it is not possible to generate an identical copy of an arbitrary quantum state. Note that the squaring of the amplitude over probabilities allows for negative and arbitrarily complex-valued amplitudes. These account for destructive interference which can be used in quantum algorithms to suppress undesired solutions. This interference is a crucial distinction of quantum algorithms relative to classical stochastic algorithms. Note also, that the need for the state vector to retain its norm (from which we inferred that operations have to be unitary) can be understood as conservation of total probability. The aforementioned basic concepts of quantum computation are direct consequences of the postulates of quantum mechanics and provide a foundation for exploring how these principles can be applied in practice through *quantum circuits*, which function as the quantum counterpart of classical algorithms.

A quantum circuit typically begins with qubits in a defined initial state, followed by the application of quantum gates that transform these states according to the intended computation. The circuit concludes with a measurement, which extracts classical information by collapsing the qubits into definite classical states. Quantum circuit diagrams serve as a useful tool for visualizing quantum algorithms. Individual qubits are repre-



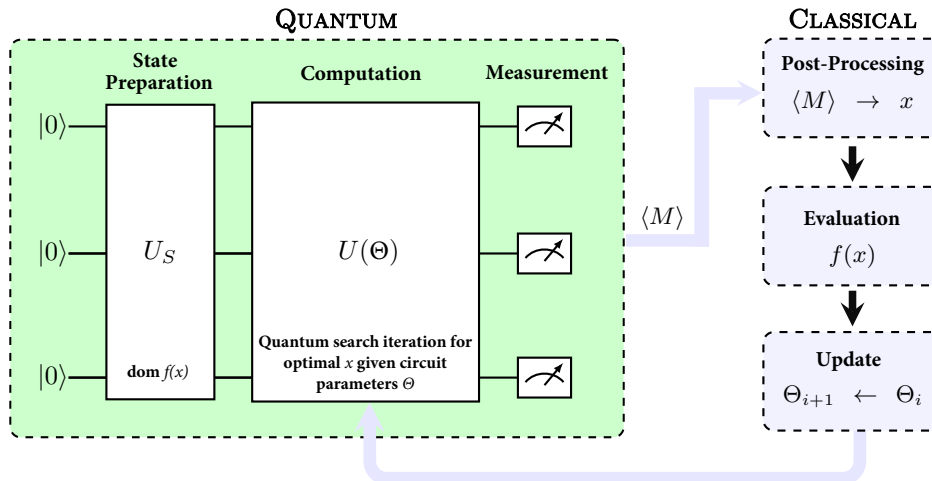
**Fig. 3** A simple quantum circuit demonstrating the combination of basic quantum gates. The circuit starts with a qubit in the  $|0\rangle$  state. The Hadamard gate ( $H$ ) is applied first, creating a superposition of  $|0\rangle$  and  $|1\rangle$ . Following this, a Pauli-X gate ( $X$ ) is applied, flipping the qubit's state. The circuit concludes with a measurement, represented by the meter symbol, which collapses the qubit's state into either  $|0\rangle$  or  $|1\rangle$ , producing a classical output.

sented as horizontal lines, and the sequence of operations, also known as gates, is indicated by their position along these lines. An example of a simple quantum circuit is depicted in Fig. 3.

Similar to classical computing there is a universal set of quantum gates, and quantum computing can simulate classical computing, while the converse is also true [57]. However, quantum computing solutions are in the complexity class of bounded-error quantum polynomial time for problems solvable in polynomial time by probabilistic quantum Turing machine with bounded error (error probability less or equal than  $1/4$ ). This class includes not only that of P but also interleaves with NP and PSPACE. Among other, that motivates research on quantum-supported problem-solving methods beyond those with an already proven significant speed-up compared to their classical counterparts such as quantum prime factorisation [160] and quantum search [138].

## 2.2 Adiabatic Quantum Computing

As an alternative, *adiabatic quantum computing* (AQC) [7] performs problem-solving by gradually evolving a quantum system toward its lowest energy state as a solution. According to the adiabatic theorem, this gradual evolution allows the system to remain in its ground state throughout the process, theoretically enabling it to perform any quantum computation. This process is designed to encode the solution to a given computational problem into this ground state, and AQC theoretically possesses the capability to perform any quantum computation, rendering it a universal approach to quantum computing. However, this theoretical potential relies on ideal conditions, including a perfectly isolated system, the ability to exert extremely precise control, and sufficiently long computation times to mitigate errors. Note that the total time for this evolution is controlled by the inverse of the energy gap, i.e., the difference between the two lowest energy eigenvalues during the system evolution, and an efficient adiabatic quantum algorithm is characterized by this gap shrinking only polynomially with problem size.



**Fig. 4** A Hybrid Quantum-Classical Optimization Workflow (adapted from [115]). The diagram illustrates a hybrid quantum-classical optimization process. The quantum part involves three main stages: **State Preparation**: An initial quantum state is prepared using the unitary operator  $U_S$ . **Computation**: The prepared state is processed by a parameterized quantum circuit  $U(\Theta)$  to search for the optimal solution  $x$  based on the parameters  $\Theta$ . **Measurement**: The quantum state is measured to obtain the expectation value  $\langle M \rangle$ . The classical part includes three steps: **Post-Processing**: The measured expectation value  $\langle M \rangle$  is processed to extract the classical variable  $x$ . **Evaluation**: The function  $f(x)$  is evaluated based on the classical variable  $x$ . **Update**: The parameters  $\Theta$  are updated using classical optimization algorithms to improve the search in the next iteration. The process iterates, with the updated parameters  $\Theta_{i+1}$  being fed back into the quantum circuit, forming a closed-loop optimization cycle between the quantum and classical computations.

A more practical application of AQC principles is *Quantum Annealing* (QA) [87], which involves guiding a quantum system towards a low-energy state but is specifically focused on finding approximate solutions to combinatorial optimization problems. Consequently, QA does not fully meet the criteria for universal quantum computing, making it more suitable for specific problem-solving scenarios rather than serving as a versatile computational tool akin to AQC.

In quantum annealing, computational problems are formulated as *Quadratic Unconstrained Binary Optimization* (QUBO) problems, mathematically described as follows:

$$\text{minimize } Q(\mathbf{x}) = \sum_{i,j} Q_{ij} x_i x_j \quad (3)$$

where  $\mathbf{x}$  is a vector of binary variables  $x_i$  (with  $x_i \in \{0, 1\}$ ), and  $Q_{ij}$  are the coefficients of the quadratic terms. The goal is to find the binary vector  $\mathbf{x}$  that minimizes the value of  $Q(\mathbf{x})$ .

In practice, while QA does not guarantee the correct solution, it remains a powerful tool for addressing classically intractable problems. Notably, the QUBO formulation in Eq. (3) is NP-complete, implying that any NP-hard problem can be mapped into a QUBO formulation with only a polynomial overhead. This suggests that rather than viewing direct QA-based quantum algorithms as merely an alternative to classical or direct gate-based quantum computing for solving

QUBO problems, it should be considered a tool that leverages quantum mechanics to provide potential advantages in specific cases, which warrant further exploration in practice.

### 2.3 Hybrid Quantum-Classical Computation

While several theoretical results demonstrate that direct (gate-based or AQC-based) quantum algorithms can solve complex problems more efficiently than classical alternatives in terms of worst-case time complexity [51, 160, 14], these results invariably assume the existence of quantum computers with a very low logical error rates. This assumption necessitates the use of *quantum error correction* to safeguard quantum information against errors and noise that may arise during computation. Quantum computers are inherently sensitive to external disturbances, hardware imperfections, and spurious decoherence, all of which can introduce errors. In contrast, NISQ devices, which represent the current state of quantum hardware, operate without full fault tolerance and have limited computational capabilities. Such machines are not yet powerful enough to outperform classical computers and it is an open challenge to find clear quantum advantage in NISQ.

Consequently, the concept of *hybrid quantum-classical computation* has been proposed to exploit near-term quantum devices and benefit from the anticipated per-

formance boost offered by quantum technologies. Specifically, in the domain of gate-based quantum computing, *variational quantum algorithms* (VQAs) have been developed to address optimization problems by utilizing both classical and quantum resources. The quantum component, often referred to as a *variational quantum circuit* (VQC) or *parameterized quantum circuit* (PQC), plays a critical role in this hybrid approach.

Within this framework, the classical input data  $x$  is initially pre-processed on a classical device to determine a normalized input quantum state for the PQC. Following this pre-processing, a VQA is executed, which comprises two sets of quantum gates:  $U_S$  and  $U(\theta)$ .

Quantum operations  $U_S$  represent the data encoding (or embedding) step and consist of a sequence of quantum gates designed to generate a quantum state that represents the classical input  $x$  as accurately as possible. The structure of  $U_S$  is task-dependent, varying according to the specific computational problem being addressed. For instance, in combinatorial optimization problems, the encoding step involves generating a quantum state that represents the classical search space of the original problem. A prominent example is the Quantum Approximate Optimization Algorithm (QAOA) [61], a well-known variational quantum algorithm for solving QUBO problems. In QAOA, the initial step of data encoding consists of generating a quantum state where the basis states encode all possible binary strings that might be solutions to the problem. Conversely, in the context of machine learning, the encoding step involves mapping the set of features into a quantum state. This mapping can be achieved through various methods, such as using the basis states, amplitude encoding, or other encoding strategies. Each of these approaches has its own advantages and disadvantages. For a more detailed discussion of these methods and their implications, see [183].

Subsequently, a sequence of parameterized quantum gates  $U(\theta)$ , referred to as the *Ansatz*, is applied with randomly initialized parameters  $\theta$ . After the execution of the PQC, which includes the quantum gates  $U_S$  and  $U(\theta)$ , the result of the measurement is classically post-processed to obtain a classical output  $f(x)$ . This output is then utilized to evaluate a task-dependent cost function. Based on this evaluation, the parameters of the PQC are updated using gradient descent or another optimization algorithm. This process is repeated iteratively in a closed loop between the classical and quantum hardware until the optimization converges or the desired performance is achieved.

The strength of this approach lies in the adaptability of the architecture, which allows for customization through learning of the gate parameters in the PQC

of the VQA to address various use cases. The entire procedure is depicted in Fig. 4.

Quantum annealing can also be utilized within a hybrid framework. In this case, the approach involves designing hybrid quantum-classical algorithms that iteratively generate large QUBO problems, which can then be solved using quantum annealing. Here, only a specific computational component is delegated to quantum annealing, while the rest of the computation remains classical.

### 3 Quantum Computing for AI

One direction of quantum AI refers to the use of quantum computing for solving computational problems in AI. As mentioned above, this concerns all subfields of AI such as automated planning, machine learning, computer vision, natural language processing, and multi-agent systems, for each of which we provide selected results and insights in this section. The other direction of QAI, namely, AI for quantum computing is covered in Section 4.

#### 3.1 Quantum Machine Learning

*Quantum machine learning* (QML) seeks to harness the principles of quantum computing to perform traditional machine learning (ML) tasks [20, 16, 17]. Although theoretical research indicates that fault-tolerant quantum computing could accelerate the training of various ML algorithms providing a computational advantage in terms of worst-case time complexity [151, 113, 114], current quantum hardware is not yet powerful enough to implement these algorithms effectively. Consequently, in QML the focus has shifted towards leveraging hybrid quantum-classical computation with VQAs (cf. Section 2.3), which aims to exploit gate-based near-term quantum devices to develop innovative models and potentially achieve performance gains from quantum technologies.

The hybrid QML approach with the use of VQAs (cf. Section 2.3) shares significant similarities with the training of classical neural networks (NNs), notably in their reliance on parameterized models, gradient-based optimization techniques, and structured layers for approximating complex functions. As such, this common ground has paved the way for the development of *quantum neural networks* (QNNs), which essentially are the use of PQCs for machine learning applications. QNNs are widely used in both supervised and reinforcement learning. In *quantum supervised learning* [112], these algorithms typically involve fitting a parameterized function

to a training dataset, that can represent complex hypothesis functions that classical models might struggle with. Therefore, the use of QNNs in supervised learning seeks to identify problem classes that are intractable for classical approaches from a learning perspective, rather than expedite the training process in terms of worst-time complexity.

Similarly, QNNs can be effectively utilized in *quantum reinforcement learning* (QRL) [122,54]. In reinforcement learning (RL), an agent interacts with an environment to maximize cumulative rewards by learning optimal policies[72]. Hybrid quantum-classical approaches of RL can employ QNNs to encode and process states and actions within the environment, leveraging a potential quantum advantage in representing and exploring complex state-action spaces. QNNs can be used to approximate value functions or policy distributions, while classical components handle the optimization of these quantum parameters through iterative updates. This hybrid strategy can enhance the learning efficiency and representational power of RL algorithms, potentially solving problems that are challenging for classical methods alone.

Recently, a quantum deep reinforcement learning method has been developed and experimentally evaluated on noisy gate-based quantum simulation for the use case of safe navigation by self-driving cars that do not need a quantum device on-board during testing but training only [163]. The QRL method leverages the model capacity of quantum neural networks, only requires a few dozen noisy qubits and relies on hybrid quantum-classical computations to be effective. The experimental evaluation provided evidence in favor of quantum utility in terms of faster and more stable training with fewer parameters compared to the classical counterpart. These results suggest that existing noisy quantum computing devices with a few tens of qubits might soon become viable alternatives to overcome the challenges faced by classical methods in, for example, enhancing autonomous systems and optimizing large distribution networks through quantum-supported reinforcement learning.

The focus on using quantum algorithms for unsupervised learning primarily involves fault-tolerant quantum approaches that can theoretically provide a computational speedup in terms of worst-time complexity. Key advancements include the introduction of quantum algorithms that utilize amplification techniques in clustering problems [5,55], the proposal of a quantum k-means and k-nearest centroid algorithms [110,88], and recent efforts focused on graph sparsification achieving by leveraging superposition access to classically stored graph weights [12].

### 3.2 Quantum Planning and Scheduling

*Quantum automated planning and scheduling* (QPS) research focuses on quantum-supported means of automated planning (QP) and scheduling (QS) in AI, and vice versa. Automated planning methods in AI can be divided into online and offline planning each with certainty or under uncertainty. Offline planning is decoupled from the subsequent execution of produced plans and gets no feedback about it during planning, while online planning is interleaved with a controlled action execution in a closed-loop manner. In many real-world applications, action planning under uncertainty is required that allows for actions with non-deterministic effects and incomplete initial states caused by only partial observability of the environment such as in partially observable Markov decision processes (POMDP). Depending on the chosen technique of planning with certainty or under uncertainty, a plan can take the form of, for example, a finite sequence of primitive actions, a conditional action plan, or an action policy on belief states with maximal expected utility. For a more comprehensive introduction to automated planning, we refer the interested reader to, for example, [66].

In any case, automated planning methods in AI are known to be computationally expensive, such as classical (state-space or planning graph-based) action planning already being PSPACE-complete, while belief states in POMDPs and non-determinism as well as interactive POMDPs make it even worse exponentially. Nevertheless, automated planning of symbolic AI is considered more explainable in general than deep learning-based POMDP solutions, and approximated online POMDP planning methods are successfully used in many practical applications such as robot navigation. The fundamental question therefore arises as to whether, to what extent and under which conditions quantum-supported AI planning, particularly for (interactive) POMDPs, is feasible and may lead to a significant reduction of planning time and space compared to the classical counterparts. As of today, not much is known in this regard yet.

For example, a classical POMDP models an agent acting in a partially observable stochastic environment. A first quantum-based POMDP model was proposed in [15], which is intended to generalize POMDPs and has the same complexity for the strategy (plan) existence problem, i.e. is PSPACE-hard for a polynomial (undecidable for infinite) time horizon. In this model, actions and observation process are represented by a super-operator on quantum-encoded environment states assumed by the agent. However, it is unclear to what extent this model is useful for quantum-supported plan-

ning and learning agents in POMDPs in concrete terms. In [44], a quantum MDP model (QMDP) is defined through a potential energy function of the quantum system under consideration. It is shown how model-based learning by approximate value iteration for POMDPs can be applied to such a model. However, how QMDPs are modeled when the energy function is unknown, and whether this model may serve as a basis for quantum-supported approximated interactive POMDP planning is not known either. Another open question is to what extent an adaptation of hybrid neuro-symbolic methods for learning-assisted planning in POMDPs [141, 46] for quantum computing would be feasible and beneficial. In [30], a quantum-supported method for optimal path planning in robot navigation is presented that leverages Grover’s quantum search in a classical tree-search procedure. It is simulated for gate-based quantum computation and shown theoretically to always provide a quantum speedup up to that of the Grover algorithm [70]. Earlier work [127] proposes a QP method that adapts and benefits from standard quantum search for planning in MDPs using Dynamic Programming (DP) and a heuristic for controlling a discrete time quantum walk on the MDP state graph, which reduces the number of states visited during a DP iteration compared to classical.

*Quantum scheduling* (QS) research actually focuses on the use of quantum computation for solving the NP-complete problem of job shop scheduling (JSS) and its modern variants that allow for more flexibility (FJSS) in the context of Industry 4.0. Roughly, the FJSS problem is to find a feasible schedule that assigns constrained job operations to multi-purpose machines with, for example, a minimal production makespan in total. There exist many different kinds of variants of this problem and classical approximate solution methods based on genetic algorithms, artificial neural networks, particle swarm optimization and reinforcement learning [191, 101]. In agent-based solutions for FJSS, individual machines, carriers, jobs or individual operations are often represented by agents in a respective multi-agent system.

However, there are only very few quantum-supported solution methods for some variants of the FJSS problem yet. These QS methods are limited to quantum genetic algorithms and quantum PSO, in which iterative evolutionary sets of suitably quantum-encoded individuals can lead to approximate optimal solutions, for example by means of quantum rotations [192, 32, 167].

In [8, 6], an iterative quantum-supported hybrid MILP (Mixed Integer Linear Programming)-based optimal solution of the JSS problem is decomposed in a way that is suitable for hybrid quantum-classical computing. The

hybrid method solves a relaxed MILP problem of the original JSS problem with a classical MILP solver to find possible optimal assignments of jobs to machines. This is followed by solving related sequencing QUBO problems that correspond to individual machines with a D-Wave quantum annealer to search for feasible schedules based on those assignments. The reported experimental results show a significant computational time speedup of this method over the classical solver Gurobi for JSS problem instances comprising of up to 280 machines and jobs in a specific scheduling dataset.

Multi-agent reinforcement learning (MARL) is one classical solution approach for FJSS with dynamic changes of optimization conditions and configurations [109, 78, 135]. Currently, there are only very few initial, feasible approaches of quantum-supported MARL (QMARL) for this specific problem class [136, 189] but without any experimental analysis and concrete insights on their benefits compared to classical counterparts yet.

Another NP-hard problem in the AI domain of automated planning and scheduling is the bin-packing problem, which is to find the minimum number of bins of fixed capacity required to pack a set of items of varying size without exceeding the bin capacities. Recently, a first QS method has been presented that solves this problem reformulated as a QUBO problem on a quantum annealer [28]. The method utilizes the Augmented Lagrangian method to account for the bin packing constraints and heuristic penalty multipliers, scales with increasing problem size but does not outperform the selected classical counterpart in runtime yet due to current limitations of quantum annealing hardware.

It is apparent that more in-depth theoretical and experimental investigations are needed on which types of computationally hard AI planning may benefit from the adoption of direct or hybrid quantum computation under which conditions and assumptions of integration compared to classical solutions in general, and for which practical use cases in particular. Research in the other direction of QPS, that is, the use of AI planning and scheduling for advanced manufacturing and operating quantum devices, is as relevant as the use of AI for the same purpose regarding classical computing devices but actually occurs less prominent in the literature (cf. Sect. 4).

### 3.3 Quantum Computer Vision

*Quantum computer vision* (QCV) research is mainly concerned with the investigation of the feasibility and benefits of quantum-supported computer vision methods for the perception of intelligent agents in AI. Re-



search in the other direction of QCV, that is, the usage of computer vision methods for advances in the building and operating quantum computing devices actually appears rather neglected.

The QCV subfield of QAI formed in the early 2000s [173] and attracted renewed attention with recent advances in quantum image processing. The latter requires the representation and processing of a given digital image on a quantum computer, and the final conversion of the processed quantum image into a classical image. Quantum image representations such as qubit lattices and normal arbitrary quantum superposition states store color image information using amplitudes, phases, or basis quantum states [107, 102]. For the processing of quantum represented images, there are already quite a few, mostly quantum annealing-based methods available for tasks such as image recognition and classification [27, 172, 89], image synthesis [49], object tracking and detection [103], graph matching [159], as well as for motion and image segmentation [13, 174, 179].

For example, the problem of unsupervised graph-based image segmentation is (a) to construct a weighted undirected graph from a given image with set of vertices (pixels), set of edges (synergies between pixels), and set of weights (similarity between pixels), and then (2) to find the best partition into disjoint subsets such that the sum of weights between different subsets is minimized. This NP-hard problem has been recently solved in [174] with a quantum-supported method that first reformulates segmentation as a graph-cut optimization problem, maps it into the topology of and runs it on the considered quantum annealer, in this case, a D-Wave Advantage, then retrieves the results of quantum measurement and eventually generates the segmentation mask from them. Despite the shared and remote access of the D-Wave device, comparative experimental evaluation results revealed that this QCV method outperformed the classical solver Gurobi on the same task in terms of runtime with only slightly sub-optimal solution quality. In this regard, it is particularly valuable when collecting labeled data is costly and speed is of essence. In [179], quantum algorithms for the same purpose are analyzed that allow to scale the number of qubits exponentially with respect to the input size to use current gate-based quantum computing devices.

Motion segmentation, on the other hand, aims to detect independent motions in two or several input images. The QCV method presented in [13] solves this problem reformulated as a QUBO problem directly on a D-Wave quantum annealer with reported on-par performance compared to classical solutions.

Recently, in [172], several variations of a quantum hy-

brid vision transformer were analyzed for solving an image classification problem in high-energy physics with the result of, again, an on-par performance compared to classical counterparts with a similar number of model parameters. For small-scale medical image datasets, it has been shown in [89] that quantum transformer models with quantum attention layers may perform better than classical vision transformers for this purpose in terms of asymptotic run time and fewer model parameters.

In [49], two quantum hybrid diffusion models for image synthesis are presented. The first model replaces convolutional ResNet layers with hybrid quantum-classical variational quantum circuits only at the vertex, while the second additionally does so in the second block of the encoder part. The experimental evaluation via simulation in PennyLane indicate that such models are of benefit in the sense that they synthesize better-quality images and converge faster with a lower number of parameters to train.

### 3.4 Quantum Natural Language Processing

*Quantum natural language processing* (QNLP) is concerned with the representation and processing of natural language through quantum computational means. The other direction of QNLP, that is the utilization of NLP means for quantum computing tasks remains fully unexplored yet. As summarized in [71], most QNLP works leverage quantum superposition to model uncertainties and ambiguity in language or entanglement to describe both composition and distribution of syntax and semantics effectively. The current QNLP methods are mainly developed for the NLP tasks of question answering, text classification and translation. In [38], Coecke and colleagues show that QNLP is not just the quantum counterpart to NLP but allows to combine linguistic structure and semantics or meanings in one quantum computational system, in fact, represent this kind of knowledge in respectively composed variational quantum circuits more efficiently and inherently than in the classical case. The common basis of most QNLP approaches is the Categorical Distributional Compositional (DisCoCat) diagram model for natural language [39] with which one can encode the meaning of words and phrases as quantum states and processes, hence as quantum circuits. In this regard, the implicit flows of meanings in respective diagrammatic reasoning due to the underlying structure are exposed rather than encoded in black boxes of neural transformer-based large language models (LLM) in classical NLP.

The potential speedup of running the DisCoCat model on quantum hardware was discussed first in [190], while

in [185, 104] the potential of quantum computing techniques such as quantum search and quantum neural networks for faster training and testing of LLMs, and vice versa, is conceptually discussed but not yet demonstrated in practice. In this context, complementary work includes quantum-inspired approaches such as the recently reported compression of LLMs with quantum-inspired tensor networks by supposedly more than 80% without compromising accuracy [123].

However, more advanced experimental insights on and development of hybrid quantum-classical QNLP methods, in particular quantum neural networks, for NLP tasks are needed but limited by current NISQ quantum computing hardware. For example, as in QCV, the potential of quantum-supported transformers compared to classical transformer neural network architectures might be worth to investigate. For more information on QNLP in Quantum AI, we refer the interested reader to, for example, [185, 71, 38]

### 3.5 Quantum Agents and Multi-Agent Systems

*Quantum multi-agent systems* (QMAS) research mainly focuses on two strains: (1) development of autonomous agents and multi-agent systems for hybrid quantum-classical computational environments, and (2) quantum-supported methods for coordination and cooperation in multi-agent systems. In modern AI, the concept of an intelligent agent and multi-agent system is at the core [156, 184] and integral part of many agent-based scientific, industrial and commercial AI applications [182]. Roughly said, in a multi-agent system, multiple homogeneous or heterogeneous agents coordinate their actions to achieve joint goals and carry out tasks flexibly, autonomously and interactively in complex environments with cooperative or competitive settings depending on the considered application problem.

There are quite some tools and frameworks for the engineering of autonomous agents and multi-agent systems for environments of classical computing [11, 19, 24, 23, 41]. However, architectures and approaches to build and operate agents with quantum computational capabilities in AI applications to run on one or multiple networked quantum computers or in hybrid quantum-classical computing environments are very rare. In fact, concepts, tools and frameworks for quantum multi-agent system programming based on the currently available quantum programming frameworks and quantum computational models are needed but still missing and topics of future research in QMAS.

Early work in this direction [92, 94] proposed a first conceptual classification and architecture of quantum multi-agent systems. More recent, a quantum modeling

approach for reactive agents with subsumption architecture has been proposed and exemplified for a simple ball picking robot in [96]. The agent state is encoded as the superposition of the tensor products of a 4-qubit perception vector with a 5-qubit action vector, and the simple quantum control circuit for the robot is shown for gate-based quantum computation but not yet evaluated.

Research on quantum-supported means of coordination and cooperation in multi-agent systems is still in its infancy. In [128], two first quantum-supported methods for coordination in a quantum multi-agent system are presented for gate-based quantum computing devices. These methods were concerned with quantum versions of Kernel-based coalition negotiation and a specific contract net protocol with embedded auction. The used quantum coalition protocol provides a marginal speedup in computation and quadratic reduction in communication between agents compared to the classical counterpart, while the quantum contract net version offers no computational speedup but more data privacy to the bidding agents.

Likewise, [31] initially discusses the potential of leveraging entangled quantum states for coordination in (mixed) multi-agent systems by means of a quantum public-goods protocol and quantum auctions for resource allocation. Their experiments on the latter indicates, for example, that the quantum version may provide more privacy than the classical (first price) auction but at the cost of lower economic efficiency.

More recently, quantum multi-agent reinforcement learning (QMARL) for adaptive coordination in collaborative settings for the quantum gate-based model [50, 136, 189] or quantum annealing [129, 100] gained some interest. For example, in [50] a QMARL version for centralized training and decentralized execution (CTDE) relies on distributed advantage actor-critic architecture with a quantum critic uniquely spread across the agents and coupling of local observation encoders through entangled input qubits over a quantum channel. That eliminates explicit sharing of local observations of agents and reduces classical communication overhead. Whereas in [136, 189] alternative QMARL versions for CTDE are proposed with a quantum critic at a central server and agents sending their local observations via a classical channel.

In any case, experimental evidence in favor of, or against some quantum utility compared to classical MARL methods is shown for simple toy domains only [50], if at all. More investigations in this regard are definitely required to assess the potential of QMARL.

One prominent class of micro-level coordination techniques for multi-agent systems [133] in competitive en-

vironments is coalition formation based on cooperative game theory. In general, coalition formation refers to situations in which groups of individually rational agents intend to work jointly in (temporary) coalitions with binding agreements in order to accomplish their tasks they cannot accomplish individually [58, 157]. Research in this area focuses on two main problems, namely, coalition negotiation and coalition structure generation [146, 145]. Roughly said, for a given coalition game  $(A, v)$  with a set  $A$  of rational agents and characteristic or coalition value function  $v$ , the goal is either (a) to negotiate a game-theoretically stable coalition structure as partition of  $A$  with individual payoffs for coalition members, or (b) to generate the optimal coalition structure that maximizes the social welfare or joint profit of the multi-agent system as a whole without individual payoff distribution. Both problems are computationally expensive, thus potential candidates for the investigation of quantum-supported solution alternatives. For coalition negotiation this holds subject to the chosen stability criteria like the exponential Kernel and Shapley value, while coalition structure generation (CSG) is a NP-hard optimization problem.

Since early work on quantum (non-)cooperative games [81, 79, 128, 29, 80, 52], only recently methods for quantum coalition structure generation have been developed and experimentally evaluated against classical counterparts revealing evidence in favor of some quantum utility [176, 177, 175]. As an example, [175] presents a direct quantum solution of the CSG problem reformulated as a QUBO problem for both quantum gate-based simulation and quantum annealing devices, while [176] proposes a hybrid quantum-classical CSG method for induced subgraph games. In particular, the latter leverages quantum annealing to iteratively solve the embedded NP-hard optimal splitting problem in only linear runtime to find the best bipartition of agents by moving down partition hierarchy until no value-increasing bipartitions remain. In fact, it explores a larger portion of the solution space compared to other approximate classical bottom-up solvers and outperforms them with its quadratic runtime in the number of agents and an expected worst-case approximation ratio of 92% on standard benchmarks.

Remarkably, quantum coalition formation methods have already been applied to practical use cases of energy management [22, 124]. The management of energy resources presents complex computational challenges, particularly with scenarios involving dynamic energy consumption and optimizing energy distribution in microgrids and large-scale power networks as well. For example, the use case problem considered in [124] is to find an optimal coalition structure of energy prosumers

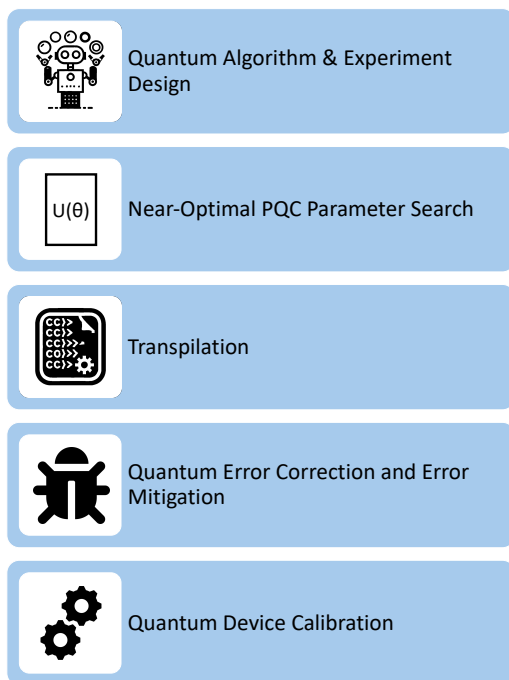
with the aim to minimize network management costs through optimized power flow analysis, that is to maximize the overall network efficiency. The results of using the quantum CSG method from [176] for solving this problem on quantum annealers showcase quantum utility compared to classical standard solutions. Advancing in this direction will facilitate the broader adoption of quantum AI solutions for real-world energy management challenges.

Another use case of quantum CSG methods in the aerospace domain was recently investigated in [178]. The problem to find constellations of large-scale low-earth satellite networks such as StarLink networks by dynamic clustering of moving LEO satellites with minimal inter- and intra-cluster communication is NP-hard. According to initial experimental results, the used quantum CSG method on a quantum annealer outperformed the classical state-of-the-art solver for this problem on real-world orbital datasets from Celestrak related to Starlink satellites.

Coalition formation methods have also been leveraged in the domain of transport and logistics for collaborative solutions of the NP-hard problem of vehicle routing and variants of the capacitive vehicle routing problem (CVRP) [116, 194, 181]. However, quantum-supported solutions for the same are still rare and topic of future research [62, 82]. For example, in [62] a CVRP is solved hybrid quantum-classical through classical clustering and optimal route planning in each cluster, while the latter is done with a direct quantum optimization algorithm for solving the corresponding TSP (Traveling Salesman Problem)-QUBO problem on a D-Wave quantum annealer. The method did not provide a clear benefit in solution quality and runtime, though the latter might change with the advent of a more advanced quantum annealer in the future.

## 4 AI for Quantum Computing

In this section, we turn our attention to the other direction of QAI, namely, the use of AI for quantum computing. It summarizes selected approaches to exploit AI, particularly ML methods, in support of the whole process from quantum algorithm and experiment design, search for near-optimal parameters, transpilation of quantum circuits, error correction during execution, and the calibration and design of quantum devices as shown in Fig. 5. It can be seen as a particular example of the field of AI for science.



**Fig. 5** Stack of tasks in the building and operation of quantum computing devices for which AI techniques (currently mainly from ML) are utilized.

#### 4.1 Quantum Algorithm and Experiment Design

Though the design and implementation of quantum algorithms in general is manually done by human experts, there are quite a few approaches to leverage ML methods in support of both the experiment design and protocol development. There is a common agreement that humans might not be predestined to design new experiments and protocols in quantum computing when recent experimental findings deliver counter-intuitive results. Most approaches to experiment design suggest the use of reinforcement learning and minimizing the influence of the human developer.

In this regard, the framework MELVIN [98] offers the use of ML methods to discover new experimental methods for creating and manipulating complex quantum states. It creates experiments with optical components, arranges them randomly, calculates and analyzes quantum states, and simplifies configurations based on user-defined criteria, if the desired properties are met. Remarkably, MELVIN produced correct and useful but unfamiliar, asymmetric techniques that are hard to understand intuitively.

Inspired by MELVIN, [119] and [180] use the physics-oriented ML approach of projective simulation based on RL for autonomous experiment design. For example, the authors in [119] applied projective simulation for developing complex photonic quantum experiments

with multi-photon states entangled in high dimensions. As a result, the approach unexpectedly rediscovered advanced experimental techniques, in particular by interacting with a simulated optical table using a set of optical elements to create experiments. It places elements on the table, analyzes the resulting quantum states, and receives rewards based on the given task and state before iterating the process.

In [180] the same approach is used to discover and optimize protocols in quantum communication over long distances. Their RL agent successfully rediscovered established protocols such as for quantum teleportation, entanglement purification, and quantum repeaters. The results suggest that RL could be effective in finding solutions that outperform human-designed ones for long-distance communication challenges, especially in scenarios with asymmetric conditions.

In contrast, [43] investigates the use of differential evolution and particle swarm optimization for improving the precision of multi-particle entanglement-free quantum phase estimation. Both methods are leveraged to fine-tune feedback policies aimed at reducing different types of noise. In fact, both methods showed superior robustness and precision compared to non-adaptive methods, particularly for scenarios where noisy, non-entangled qubits serve as sensors in quantum sensing and metrology.

Related research on explainable AI (XAI) for quantum computing is concerned with the use of XAI techniques not only to explain the designed quantum algorithms but quantum circuits in general and PQCs for QML (cf. Section 3.1) in particular. For example, [165] reports that PQCs introduce probabilistic errors due to quantum measurements, which complicates the use of traditional XAI methods. Moreover, with the phase space of a quantum circuit expanding exponentially with the number of qubits, executing XAI methods in polynomial time becomes challenging. The authors evaluated the adaptation of the XAI methods IG (Integrated Gradients) and baseline SHAP (SHapley Additive exPlanations) for PQCs using truncated Fourier series. Based on their findings, the authors introduced qSHAP (quantum SHAP), which scales with the number of features rather than qubits, making it suitable for larger quantum circuits, and shows more robustness against noise compared to the other two approaches.

In [77], SHAP is investigated in the form of quantum Shapley (QShap) values to assess the impact of single or group of quantum gates in PQCs, similar to evaluating feature importance in classical ML. QShap values are model-agnostic within quantum domains, treating gates as players in a coalition game to measure their contribution to tasks like expressibility, entanglement

capability, or classification quality (see also [161]). To handle the uncertainty in quantum computing, QShap values are based on uncertain Shapley values, which account for measurement noise and decoherence. The authors tested QShap values on quantum support vector machines (QSVM), quantum generative adversarial networks (QGANs) and QAOA (cf. Section 2.1), reporting on the significance of individual gates in these algorithms.

In contrast, using the inception or so-called deep dreaming method, [85] analyzes a neural networks' understanding of quantum optics experiments. This method shows the ability of the used neural network to find novel configurations beyond the initial data and reveals a progression from simple to complex feature recognition in the layers of the network.

Overall, most progress is made in the area of experiment design with algorithm design lacking behind. As to XAI, first results, primarily based on Shapley values, were presented. Further, XAI will get extremely important in QML as the EU AI Act defines transparency rules, besides others, for all applications of AI, including QAI. Further, a better understanding of the influence of individual gates, as proposed in [77], might help in optimizing the depth and width of PQCs during the transpilation process.

#### 4.2 Near-Optimal PQC Parameter Search

*Warm-starting* of quantum algorithms refers to the process of finding near-optimal parameters of PQCs in hybrid quantum-classical algorithms without executing them. They aim to reduce the number of quantum circuit evaluations required during the optimization. Classical warm-starting protocols rely on relaxations of the problem which have to be solved classically to find a good initial state [56]. In the following, we summarize selected approaches that involve classical ML algorithms for this purpose.

For example, [9] investigates four regression-based machine learning methods to speedup the QAOA optimization loop. These methods include Gaussian process regression, linear model, regularized support vector machine, and regression tree. Of these, Gaussian process regression achieved the best results. In particular, despite using a relatively small training set, this approach generalized well and led to an average reduction of 44.9% in the number of optimization iterations across all local optimization procedures.

In [84] and [105], the integration of graph neural networks (GNNs) with QAOA is explored to improve the parameter initialization. Various GNN architectures such

as graph convolutional networks, graph attention networks, graph isomorphism networks, and GraphSAGE [75] were investigated in [105]. Experimental evaluation with benchmarks for their performance in initializing QAOA showed that all GNNs provided a more stable and reliable initialization compared to random initialization. One of the strengths of GNN-based initialization is its ability to generalize across different graph sizes. This means that a GNN trained on smaller graph instances can still perform well on larger instances, providing a flexible and efficient warm-starting mechanism for QAOA. Including graphs of different sizes the GNN method in [84] generated solutions of approximately 95% of the quality of the Goemans-Williamson algorithm.

The ability of matrix product states (MPS) for improving QML algorithms is explored in [48] by addressing the issue of vanishing gradients, or barren plateaus that complicate the training of PQCs. The authors propose optimizing MPS with classical methods such as density matrix renormalization group and time-evolving block decimation to approximate solutions for quantum circuits. The optimized MPS is converted into unitary matrices for quantum circuits using decompositions. Experiments showed that MPS-initialized circuits for problems like Max-Cut on a six-vertex graph and image classification with Fashion-MNIST achieved better performance, that is, with fewer gradient steps and faster convergence compared to random or identity matrix initialization.

Overall, ML algorithms show promise in the warm-starting of variational quantum algorithms. Future research might aim at methods to find the optimal parameters of the PQCs of such algorithms directly using some sort of ML algorithms and therefore replacing the classical optimization loop completely. In a similar direction, future ML-based algorithms might also assist with choosing the best hyper-parameters such as for the Ansatz circuit and the number of layers, which was not explicitly mentioned by the papers referenced in this section.

#### 4.3 Transpilation of Quantum Circuits

Given the challenges of NISQ devices, such as low fidelity and short coherence times, quantum circuits must be designed with minimal gates. Afterwards, quantum circuit transpilation includes the decomposition of non-native gates into native ones and the addition of SWAP-gates for connectivity compatibility, followed by an optimization to minimize resource usage. IBM Qiskit [86] divides this quantum circuit transpilation process into

six stages: Init (prepares the circuit by unrolling instructions and validating them), layout (maps virtual to physical qubits considering connectivity and calibration), routing (ensures backend compatibility with additional gates), translation (converts gates to the backend’s basis set), optimization (reduces circuit complexity), and scheduling (applies hardware-aware scheduling). Currently, ML-based algorithms are utilized for most of the steps to find optimal transpiled circuits. These include evolutionary algorithms, deep neural networks, and RL methods.

*Evolutionary algorithms* including genetic algorithms, genetic programming, ant colony optimization, and evolutionary deep neural networks are used in all stages of the transpilation process [99]. Further, [45] and [60] utilize evolutionary algorithms such as evolutionary deep learning for an initial qubit mapping.

*Neural networks* are primarily used for circuit mapping and circuit optimization. In [1], neural networks are utilized to significantly speed up the circuit mapping process while maintaining mapping accuracy and reducing the required computational resources. The experiments showed this on a 5-qubit IBM Q processors and compared to classical state-of-the-art mapping algorithms over a special data set for training [2]. However, [45] argue that the method and dataset above lacks correctness, generalization, and diversity of the dataset.

In [121], on the other hand, a framework based on long-short term memory (LSTM) networks is proposed to decide whether a quantum circuit can be optimized in the first place. The reasoning is that, as mentioned above, the optimization process is a complex task and it is not guaranteed to receive a more resource efficient quantum circuit. Thus, if their framework predicts that the quantum circuit cannot be optimized, classical computing resources are saved. Results show that their ML model is able to decide if an arbitrary quantum circuit can be optimized above a certain threshold with an accuracy of 96.7%.

Similar to [121], Quetschlich et al. introduce in [143] and [144] an algorithm, which finds the optimal compilation flow for a given quantum circuit. [144] builds on their previous work [143] by shifting from reinforcement learning to a simpler, scalable supervised learning approach for predicting optimal compilation options. The proposed tool recommends the best compiler options for quantum circuits, targeting end-users who may struggle with option selection. Using a statistical classifier, the tool predicts the best technology, device and compiler settings with accurate results for about 75% of unseen test circuits.

Paler et al. [134] introduce a quantum circuit mapping heuristic, named QXX, and its ML-enhanced version

QXX-MLP to improve layout, routing, and optimization stages. The goal is to map circuit qubits to physical qubits while optimizing the circuit depth. QXX-MLP uses a multi-layer perceptron to infer optimal parameter values for QXX, reducing circuit depth by a Gaussian function estimation. QXX achieves depth ratios about 30% lower than Qiskit on shallow circuits and performs competitively with Qiskit and TKET on deeper circuits. Beyond that QXX-MLP achieves almost instantaneous layout performance, significantly reducing the quantum circuit compilation time.

*Reinforcement learning* is primarily utilized to find near-optimal gate synthesis for a specific hardware backend. A deep reinforcement learning method for approximating single-qubit unitaries is proposed in [126]. This method aims to reduce the overall execution time by learning a general strategy through a single pre-compilation procedure. The authors highlight the trade-off between the length of the sequence and execution time, suggesting that their approach can potentially allow for real-time operations by improving on this trade-off.

In [60], the problem of quantum circuit placement [118] is addressed as a bilevel optimization problem to minimize SWAP counts. The authors use a deep reinforcement learning algorithm for the lower-level optimization, improving SWAP costs through state space encoding. For the upper level, as mentioned earlier, an evolutionary algorithm is used. This ML-based framework reduces SWAP gates by up to 100% and runtime costs by up to 40 times compared to heuristic methods. Fosel et al. [63] describe a deep RL agent that optimizes quantum circuits tailored to specific hardware, improving efficiency for near-term devices. Their method reduced average circuit depth by 27% and gate count by 15% in 12-qubit random circuits, demonstrating effective resource reduction. The RL agent also scales well, optimizing larger circuits effectively.

Recently, Kremer et al. [97] presented an RL-based method for synthesizing quantum circuits, including Clifford, Linear Function, and Permutation circuits, which directly matches native device instructions and constraints. This approach eliminates the need for extra transpilation steps, optimizing the transpiling process. It also enhances circuit routing, reducing the two-qubit gate depths and counts more effectively than heuristics like SABRE. That led to a better performance on quantum devices with up to 133 qubits. For 8- to 10-qubit quantum volume [137] circuits, the RL method achieved a reduction around 20% in the CNOT depth.

In contrast to the aforementioned ML algorithms, the authors of [171] discuss the effectiveness of classical automated reasoning and formal methods in AI, namely, decision diagrams, SAT solvers, and graphical calculus-

based methods such as the ZX-calculus, for the compilation of quantum circuits. The considered compilation tasks are classical simulation, optimization, synthesis, and equivalence checking of quantum circuits. As one result, the authors expect that the use of automated reasoning methods may play a role in other quantum computing applications as well, such as for the finding of ground states, phase transitions, and quantum error correction.

Overall, we have seen several ML algorithms targeting one or more transpilation stages. We are currently at a stage, where first ML-enhanced transpilers are rolled out to the end-users for testing in IBM Qiskit [97]. However, as long as we have only access to NISQ-devices we are always striving to find shallower circuits. Furthermore, besides reducing the depth of quantum circuits, also reducing the width, that is, the number of qubits required, is another topic of future research in this context.

#### 4.4 Quantum Error Correction and Mitigation

*Quantum error correction* and *error mitigation* are both relevant processes on our way to fault-tolerant quantum computing. While error correction, as the name suggests, actively corrects errors during the execution of a quantum circuit, error mitigation targets the readout error and therefore is a classical post-processing step.

A significant challenge with quantum error mitigation is the necessity for a large number of circuits, which have to be run in advance. In essence, each possible bitstring has to be tested for read-out errors. In [106], ML is used to improve quantum error mitigation by predicting near noise-free values from noisy quantum output. Their key innovation is that trained ML models can mitigate errors without additional circuits, reducing overhead compared to zero-noise extrapolation. The authors explore several models, including linear regression, random forests, multi-layer perceptrons, and graph neural networks, finding that all except graph neural networks outperform traditional methods, with random forests consistently performing best. In fact, ML-based quantum error mitigation reduces the quantum resource overhead by 30% and the runtime overhead by 50% compared to zero-noise extrapolation.

Quantum error correction is actively researched as it will deliver fault-tolerant quantum computing. In [42], an ML algorithm for continuous quantum error correction is proposed. The approach facilitates recurrent neural networks to identify bit-flip errors in continuous noisy syndrome measurements.

Further, in [74] a deep neural network decoder for quantum error correction on IBM quantum processors was

developed and benchmarked. The study demonstrates the DNN decoder's capability to efficiently process syndrome data and correct errors, outperforming the traditional Minimum-Weight-Perfect-Matching method in certain aspects. The performance of the DNN decoder was validated through simulations and experiments on IBM devices, showing promise for real-time, scalable error correction, which is crucial for fault-tolerant quantum computing.

Overall, the research on ML algorithms for quantum error correction and quantum error mitigation is still in its early stages. In addition, the use of AI methods other than from ML for quantum error correction remains to be investigated.

#### 4.5 Calibration of Quantum Computing Devices

The *calibration and design of quantum computing devices* such as superconducting or ion trap gate-based quantum computers, quantum annealers and quantum sensors can benefit from classical ML. This is originally an application of the field of quantum optimal control theory [67, 95], which aims at finding the right analogue controls for given hardware to perform desired quantum tasks. This field has already used mathematical tools similar to those of AI and is now rephrasing a lot of its work in modern AI language [186, 154, 188]. In fact, this has recently led to the tongue-in-cheek remark statement, that AI is in fact a subfield of quantum control theory.

These methods typically have a model-based and an experiment-based part and ideally also use experimental data to improve the model, in order to create a digital twin and provide a shortcut for optimal control. This is done in isolation e.g. in [166, 40, 76, 69, 164, 65] and integrated in the commercial offerings of a number of startups [155, 18, 125].

For example, [131] presents a practical, efficient, and model-independent ML method for Bayesian parameter estimation (BPE) in quantum systems. Traditional BPE methods often require extensive calibration and explicit modeling of the measurement apparatus, which makes them impractical for complex systems. The authors frame parameter estimation as a classification problem solved using supervised learning techniques, where the output of the neural network is a Bayesian posterior distribution centered at the true parameter value, bounded by Fisher information. This approach requires fewer calibration measurements and is model-independent outperforming conventional calibration-based BPE. In [187], the energy spectrum of a Hamiltonian on a superconducting quantum device is predicted, outperforming the current state-of-the-art by over 20%. The

method uses multi-target regression to predict multiple related variables, uncovering relationships between them by employing explainable AI techniques. This approach significantly improves the accuracy of quantum device calibration.

To simplify the calibration process, [35] introduces a method to calibrate quantum photonic sensors using neural networks. The approach relies on data collected using available probe states, reducing overhead, and implicitly accounting for imperfections. The neural network demonstrated robustness to noise and scalability, making it suitable for future quantum technologies. Quantum hardware design can be made more systematic and be improved by AI as well, as shown, for example, in [64,120].

Overall, the calibration of quantum devices can be practically improved through the use of ML. In conjunction with ML-assisted experiment design (see Sec. 4.1) one even might find advantages in the use of ML for designing new quantum computers with better fidelities and connectivity.

## 5 Conclusions

Research in the interdisciplinary and nascent field of quantum AI is concerned with the use of quantum computing for addressing computationally hard problems in AI, and vice versa. So far, an impressive progress was made in both directions of QAI research. In fact, there are quite a few quantum-supported solutions of selected hard optimization problems in AI with different degrees of potential quantum utility for relevant use cases in various domains such as manufacturing, automated driving, finance, and energy management. In addition, initial research revealed that quantum computing itself may benefit from the use of ML for optimizing the control, performance and calibration of quantum computational devices.

In our view, future QAI research across all its subfields should focus even more on investigations under what conditions and settings in concrete (industrial) use cases are direct or hybrid quantum-classical solutions feasible with what quantum utility in practice. Further, the appropriate and timely transition to and investigation of the feasibility and potential of nowadays QAI methods on future built non-NISQ devices is another challenge. That is particularly important in the context of current expectations of the economic value of QAI applications in relevant industries. Among other, this requires both the physics and computer science communities to even more join forces, and on the other hand, a further, sustainable support of research

on both quantum AI and the building of more resourceful, fault-tolerant quantum computing devices at government and industry level worldwide.

**Acknowledgements** We gratefully acknowledge support of this work by the German Ministry for Education and Research (BMB+F) under project grant Q(AI)2.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. G. Acampora and R. Schiattarella. Deep neural networks for quantum circuit mapping. *Neural Computing and Applications*, 33(20):13723–13743, 2021.
2. G. Acampora, R. Schiattarella, and A. Troiano. A dataset for quantum circuit mapping. *Data in Brief*, 39:107526, 2021.
3. D. Aharonov. Quantum computation. *Annual Reviews of Computational Physics VI*, pages 259–346, 1999.
4. D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Journal on Computing*, 37(1):166–194, 2007.
5. E. Aïmeur, G. Brassard, and S. Gambs. Quantum speed-up for unsupervised learning. *Machine Learning*, 90:261–287, 2013.
6. A. Ajagekar, K. A. Hamoud, and F. You. Hybrid classical-quantum optimization techniques for solving mixed-integer programming problems in production scheduling. *IEEE Transactions on Quantum Engineering*, 3:1–16, 2022.
7. A. Ajagekar, T. Humble, and F. You. Quantum computing based hybrid solution strategies for large-scale discrete-continuous optimization problems. *Computers & Chemical Engineering*, 132:106630, 2020.



8. A. Ajagekar, T. Humble, and F. You. Quantum computing based hybrid solution strategies for large-scale discrete-continuous optimization problems. *Computers & Chemical Engineering*, 132:106630, 2020.
9. M. Alam, A. Ash-Saki, and S. Ghosh. Accelerating quantum approximate optimization algorithm using machine learning. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 686–689, 2020.
10. Y. Alexeev et al. Quantum-centric supercomputing for materials science: A perspective on challenges and future directions. *Future Generation Computer Systems*, 160:666–710, 2024.
11. A. Antakli et al. Ajan: An engineering framework for semantic web-enabled agents and multi-agent systems. In P. Mathieu, F. Dignum, P. Novais, and F. De la Prieta, editors, *Advances in Practical Applications of Agents, Multi-Agent Systems, and Cognitive Mimetics. The PAAMS Collection*, pages 15–27, Cham, 2023. Springer Nature Switzerland.
12. S. Apers and R. de Wolf. Quantum speedup for graph sparsification, cut approximation and laplacian solving. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 637–648, 2020.
13. F. Arrigoni, W. Menapace, M. S. Benkner, E. Ricci, and V. Golyanik. Quantum motion segmentation. In *European Conference on Computer Vision*, pages 506–523. Springer, 2022.
14. F. Arute, , et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
15. J. Barry, D. T. Barry, and S. Aaronson. Quantum partially observable markov decision processes. *Phys. Rev. A*, 90:032311, Sep 2014.
16. C. Bauckhage et al. Quantum machine learning: State of the art and future directions.
17. C. Bauckhage et al. Quantum machine learning, eine analyse zu kompetenz, forschung und anwendung. *Fraunhofer IAIS*, 2020.
18. Y. Baum et al. Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer. *PRX Quantum*, 2:040324, Nov 2021.
19. F. Bergenti et al. *Developing Agent-Based Applications with JADE*, pages 191–214. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
20. J. Biamonte et al. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.
21. G. Birkhoff and J. V. Neumann. The logic of quantum mechanics. *Annals of Mathematics*, 37(4):823–843, 1936.
22. J. Bleninger et al. Quantum optimization for the future energy grid: Summary and quantum utility prospects, 2024.
23. R. H. Bordini, J. F. Hübner, and M. Wooldridge. *Programming multi-agent systems in AgentSpeak using Jason*, volume 15. John Wiley & Sons, 2007.
24. L. Braubach et al. Jadex: A bdi-agent system combining middleware and reasoning. In *Software Agent-Based Applications, Platforms and Development Kits*, pages 143–168, Basel, 2005. Birkhäuser Basel.
25. O. Burkacky, L. Pautasso, and N. Mohr. Will quantum computing drive the automotive future. *Mckinsey & Company*, 1:33–38, 2020.
26. D. Castelvecchi. Quantum internet milestone takes entanglement out of the lab and into cities. *Scientific American*, 2024.
27. G. Cavallaro et al. Approaching remote sensing image classification with ensembles of support vector machines on the d-wave quantum annealer. In *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, pages 1973–1976, 2020.
28. L. Cellini, A. Macaluso, and M. Lombardi. Qal-bp: an augmented lagrangian quantum approach for bin packing. *Scientific Reports*, 14(1):5142, 2024.
29. J. Chappell et al. N-player quantum games in an epr setting. *PLoS One*, 7(5):e36404, 2012.
30. A. Chella, S. Gaglio, G. Pilato, F. Vella, and S. Zammito. A quantum planner for robot motion. *Mathematics*, 10(14):2475, 2022.
31. K.-Y. Chen, T. Hogg, and B. A. Huberman. Behavior of multi-agent protocols using quantum entanglement. In *AAAI Spring Symposium: Quantum Interaction*, pages 1–8, 2007.
32. W. Chen, H. Yang, and Y. Hao. Scheduling of dynamic multi-objective flexible enterprise job-shop problem based on hybrid qpso. *IEEE Access*, 7:127090–127097, 2019.
33. Y. A. Chen et al. An integrated space-to-ground quantum communication network over 4,600 kilometres. *Nature*, 589(7841):214–219, 2021.
34. M. L. D. Chiara et al. *Quantum Computational Logics: A Survey*, pages 229–271. Springer Netherlands, Dordrecht, 2003.
35. V. Cimini, I. Gianani, N. Spagnolo, F. Leccese, F. Sciarrino, and M. Barbieri. Calibration of quantum sensors by neural networks. *Phys. Rev. Lett.*, 123:230502, Dec 2019.
36. V. Cimini, M. Valeri, E. Polino, S. Piacentini, F. Ceccarelli, G. Corrielli, N. Spagnolo, R. Osellame, and F. Sciarrino. Deep reinforcement learning for quantum multiparameter estimation. *Advanced Photonics*, 5(1):016005, 2023.
37. L. Clinton, T. Cubitt, B. Flynn, F. M. Gambetta, J. Klassen, A. Montanaro, S. Piddock, R. A. Santos, and E. Sheridan. Towards near-term quantum simulation of materials. *Nature Communications*, 15(1):211, 2024.
38. B. Coecke, G. de Felice, K. Meichanetzidis, and A. Toumi. Foundations for near-term quantum natural language processing. *arXiv preprint arXiv:2012.03755*, 2020.
39. B. Coecke, M. Sadrzadeh, and S. Clark. Mathematical foundations for a compositional distributional model of meaning. *arXiv preprint arXiv:1003.4394*, 2010.
40. J. H. Cole et al. Identifying an experimental two-state hamiltonian to arbitrary accuracy. *Phys. Rev. A*, 71:062312, Jun 2005.
41. R. W. Collier et al. Reflecting on agent programming with agentspeak (1). In *PRIMA 2015: Principles and Practice of Multi-Agent Systems: 18th International Conference, Bertinoro, Italy, October 26-30, 2015, Proceedings 13*, pages 351–366. Springer, 2015.
42. I. Convy et al. Machine learning for continuous quantum error correction on superconducting qubits. *New Journal of Physics*, 24(6):063019, 2022.
43. N. F. Costa et al. Benchmarking machine learning algorithms for adaptive quantum phase estimation with noisy intermediate-scale quantum sensors. *EPJ Quantum Technology*, 8(1):16, 2021.
44. D. Crawford et al. Reinforcement learning using quantum boltzmann machines. *Quantum Info. Comput.*, 18(1–2):51–74, feb 2018.

45. Z. A. Dahi et al. An evolutionary deep learning approach for efficient quantum algorithms transpilation. In S. Smith, J. Correia, and C. Cintrano, editors, *Applications of Evolutionary Computation*, pages 240–255, Cham, 2024. Springer Nature Switzerland.
46. M. H. Danesh, P. Cai, and D. Hsu. LEADER: Learning attention over driving behaviors for planning under uncertainty. In *6th Annual Conference on Robot Learning*, 2022.
47. J. Dargan. How close are we to quantum artificial intelligence? *The Quantum Insider*, 2023.
48. J. Dborin, F. Barratt, V. Wimalaweera, L. Wright, and A. G. Green. Matrix product state pre-training for quantum machine learning. *Quantum Science and Technology*, 7(3):035014, may 2022.
49. F. De Falco, A. Ceschini, A. Sebastianelli, B. Le Saux, and M. Panella. Quantum hybrid diffusion models for image synthesis. *KI-Künstliche Intelligenz*, pages 1–16, 2024.
50. A. DeRieux and W. Saad. eqmarl: Entangled quantum multi-agent reinforcement learning for distributed cooperation over quantum channels, 2024.
51. D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.
52. I. Dey et al. Quantum game theory meets quantum networks. *Wireless Commun.*, 31(4):90–96, aug 2024.
53. C. Dilmegani. In-depth guide to quantum artificial intelligence in 2022. *AI Multiple*, 2022.
54. V. Dunjko et al. Advances in quantum reinforcement learning. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, page 282–287. IEEE Press, 2017.
55. C. Durr and P. Hoyer. A Quantum algorithm for finding the minimum. *arXiv:quant-ph/9607014*, 7 1996.
56. D. J. Egger et al. Warm-starting quantum optimization. *Quantum*, 5:479, 2021.
57. J. Eisert and M. M. Wolf. *Quantum Computing*, pages 253–286. Springer US, Boston, MA, 2006.
58. E. Elkind et al. Computational coalition formation. *Multiagent systems*, pages 329–380, 2013.
59. European Quantum Internet Alliance. <https://quantuminternetalliance.org/>.
60. H. Fan et al. Optimizing quantum circuit placement via machine learning. In *Proceedings of the 59th ACM/IEEE Design Automation Conference, DAC '22*, page 19–24, New York, NY, USA, 2022. Association for Computing Machinery.
61. E. Farhi and A. W. Harrow. Quantum supremacy through the quantum approximate optimization algorithm, 2019.
62. S. Feld et al. A hybrid solution method for the capacitated vehicle routing problem using a quantum annealer. *Frontiers in ICT*, 6, 2019.
63. T. Fösel et al. Quantum circuit optimization with deep reinforcement learning, 2021.
64. P. García-Azorín et al. Robust multi-mode superconducting qubit designed with evolutionary algorithms, 2024.
65. V. Gebhart et al. Learning quantum systems. *Nature Reviews Physics*, 5(3):141–156, Mar. 2023.
66. M. Ghallab, D. Nau, and P. Traverso. *Automated planning and acting*. Cambridge University Press, 2016.
67. S. J. Glaser et al. Training schrödinger’s cat: Quantum optimal control: Strategic report on current status, visions and goals for research in europe. *The European Physical Journal D*, 69:1–24, 2015.
68. P. Gohel and M. Joshi. Quantum time series forecasting. In *Sixteenth International Conference on Machine Vision (ICMV 2023)*, volume 13072, pages 390–398. SPIE, 2024.
69. C. Granade, C. Ferrie, I. Hincks, S. Casagrande, T. Alexander, J. Gross, M. Kononenko, and Y. Sanders. Qinfer: Statistical inference software for quantum applications. *Quantum*, 1:5, 2017.
70. L. K. Grover. 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, 1996. Association for Computing Machinery.
71. R. Guarasci et al. Quantum natural language processing: Challenges and opportunities. *Applied Sciences*, 12(11), 2022.
72. D. Gupta and M. Klusch. Hylear: Hybrid deep reinforcement learning and planning for safe and comfortable automated driving. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, pages 1–8, 2023.
73. G. Görz, U. Schmid, and T. Braun, editors. *Handbuch der Künstlichen Intelligenz*. De Gruyter Oldenbourg, Berlin, Boston, 2021.
74. B. Hall et al. Artificial neural network syndrome decoding on ibm quantum processors. *Phys. Rev. Res.*, 6:L032004, Jul 2024.
75. W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 1025–1035, Red Hook, NY, USA, 2017. Curran Associates Inc.
76. R. Harper, S. T. Flammia, and J. J. Wallman. Efficient learning of quantum noise. *Nature Physics*, 16(12):1184–1188, 2020.
77. R. Heese, T. Gerlach, S. Mücke, S. Müller, M. Jakobs, and N. Piatkowski. Explaining quantum circuits with shapley values: Towards explainable quantum machine learning, 2023.
78. F. Inal Ali et al. A multi-agent reinforcement learning approach to the dynamic job shop scheduling problem. *Sustainability*, 15(10), 2023.
79. A. Iqbal. Studies in the theory of quantum games, 2006.
80. A. Iqbal, J. M. Chappell, C. Szabo, and D. Abbott. Resolving game theoretical dilemmas with quantum states. *Quantum Information Processing*, 23(1):5, 2023.
81. A. Iqbal and A. Toor. Quantum cooperative games. *Physics Letters A*, 293(3-4):103–108, 2002.
82. H. Irie et al. Quantum annealing of vehicle routing problem with time, state and capacity. In *Quantum Technology and Optimization Problems: First International Workshop, QTOP 2019, Munich, Germany, March 18, 2019, Proceedings 1*, pages 145–156. Springer, 2019.
83. B. Jaderberg, A. A. Gentile, A. Ghosh, V. E. Elfving, C. Jones, D. Vodola, J. Manobianco, and H. Weiss. Potential of quantum scientific machine learning applied to weather modelling. *arXiv preprint arXiv:2404.08737*, 2024.
84. N. Jain et al. Graph neural network initialisation of quantum approximate optimisation. *Quantum*, 6:861, 2022.
85. T. Jaouni, S. Arlt, C. Ruiz-Gonzalez, E. Karimi, X. Gu, and M. Krenn. Deep quantum graph dreaming: deciphering neural network insights into quantum experiments. *Machine Learning: Science and Technology*, 5(1):015029, 2024.
86. A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S.

- Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta. Quantum computing with qiskit, 2024.
87. T. Kadowaki and H. Nishimori. Quantum annealing in the transverse ising model. *Phys. Rev. E*, 58:5355–5363, Nov 1998.
  88. I. Kerenidis, J. Landman, A. Luongo, and A. Prakash. q-means: A quantum algorithm for unsupervised machine learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
  89. I. Kerenidis, N. Mathur, J. Landman, M. Strahm, Y. Y. Li, et al. Quantum vision transformers. *Quantum*, 8:1265, 2024.
  90. C. M. Kerskens and D. L. Pérez. Experimental indications of non-classical brain functions. *Journal of Physics Communications*, 6(10):105001, oct 2022.
  91. M. H. Khatami, U. C. Mendes, N. Wiebe, and P. M. Kim. Gate-based quantum computing for protein design. *PLOS Computational Biology*, 19(4):e1011033, 2023.
  92. M. Klusch. Toward quantum computational agents. In M. Nickles, M. Rovatsos, and G. Weiss, editors, *Agents and Computational Autonomy*, pages 170–186, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
  93. M. Klusch, J. Lässig, and F. K. Wilhelm. Quantum technologies and ai – interview with tommaso calarco, 2024.
  94. M. Klusch and R. Schubotz. Programming and simulation of quantum search agents. In *2007 IEEE International Conference on Systems, Man and Cybernetics*, pages 246–252, 2007.
  95. C. P. Koch et al. Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in europe. *EPJ Quantum Technology*, 9(1):19, 2022.
  96. A. Koukam et al. Towards a quantum modeling approach to reactive agents. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 130–136, 2021.
  97. D. Kremer, V. Villar, H. Paik, I. Duran, I. Faro, and J. Cruz-Benito. Practical and efficient quantum circuit synthesis and transpiling with reinforcement learning, 2024.
  98. M. Krenn, M. Malik, R. Fickler, R. Lapkiewicz, and A. Zeilinger. Automated search for new quantum experiments. *Phys. Rev. Lett.*, 116:090405, Mar 2016.
  99. J. Kusyk et al. Survey on quantum circuit compilation for noisy intermediate-scale quantum computers: Artificial intelligence to heuristics. *IEEE Transactions on Quantum Engineering*, 2:1–16, 2021.
  100. M. Kölle, F. Topp, T. Phan, P. Altmann, J. Nüßlein, and C. Linnhoff-Popien. Multi-agent quantum reinforcement learning using evolutionary optimization, 2024.
  101. M. E. Leusin et al. Solving the job-shop scheduling problem in the industry 4.0 era. *Technologies*, 6(4), 2018.
  102. H.-S. Li, S. Song, P. Fan, H. Peng, H. ying Xia, and Y. Liang. Quantum vision representations and multi-dimensional quantum transforms. *Information Sciences*, 502:42–58, 2019.
  103. J. Li and S. Ghosh. Quantum-soft qubo suppression for accurate object detection. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *Computer Vision – ECCV 2020*, pages 158–173, Cham, 2020. Springer International Publishing.
  104. Z. Liang, J. Cheng, R. Yang, H. Ren, Z. Song, D. Wu, X. Qian, T. Li, and Y. Shi. Unleashing the potential of llms for quantum computing: A study in quantum architecture design. *arXiv preprint arXiv:2307.08191*, 2023.
  105. Z. Liang, G. Liu, Z. Liu, J. Cheng, T. Hao, K. Liu, H. Ren, Z. Song, J. Liu, F. Ye, et al. Graph learning for parameter prediction of quantum approximate optimization algorithm. *arXiv preprint arXiv:2403.03310*, 2024.
  106. H. Liao, D. S. Wang, I. Sitdikov, C. Salcedo, A. Seif, and Z. K. Mineev. Machine learning for practical quantum error mitigation. *arXiv preprint arXiv:2309.17368*, 2023.
  107. M. Lisnichenko and S. Protasov. Quantum image representation: A review. *Quantum Machine Intelligence*, 5(1):2, 2023.
  108. A. Litt, C. Eliasmith, F. W. Kroon, S. Weinstein, and P. Thagard. Is the brain a quantum computer? *Cognitive Science*, 30(3):593–603, 2006.
  109. R. Liu, R. Piplani, and C. Toro. A deep multi-agent reinforcement learning approach to solve dynamic job shop scheduling problem. *Computers & Operations Research*, 159:106294, 2023.
  110. S. Lloyd, M. Mohseni, and P. Rebentrost. Quantum algorithms for supervised and unsupervised machine learning, 2013.
  111. H. Ma, J. Liu, H. Shang, Y. Fan, Z. Li, and J. Yang. Multiscale quantum algorithms for quantum chemistry. *Chemical Science*, 14(12):3190–3205, 2023.
  112. A. Macaluso. Quantum supervised learning. *KI-Künstliche Intelligenz*, pages 1–15, 2024.
  113. A. Macaluso, L. Clissa, S. Lodi, and C. Sartori. Quantum splines for non-linear approximations. In *Proceedings of the 17th ACM International Conference on Computing Frontiers*, pages 249–252, 2020.
  114. A. Macaluso, M. Klusch, S. Lodi, and C. Sartori. Maqa: a quantum framework for supervised learning. *Quantum Information Processing*, 22(3):159, 2023.
  115. A. Macaluso, F. Orazi, M. Klusch, S. Lodi, and C. Sartori. A variational algorithm for quantum single layer perceptron. In *International Conference on Machine Learning, Optimization, and Data Science*, pages 341–356. Springer, 2022.
  116. S. Mak, L. Xu, T. Pearce, M. Ostroumov, and A. Brintrup. Fair collaborative vehicle routing: A deep multi-agent reinforcement learning approach. *Transportation Research Part C: Emerging Technologies*, 157:104376, 2023.
  117. J. Marti-Guerrero et al. Quantum artificial intelligence: A tutorial. In *Proceedings of the 31st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2023)*, October 4-6 2023.
  118. D. Maslov et al. Quantum circuit placement: optimizing qubit-to-qubit interactions through mapping quantum circuits into a physical experiment. In *Proceedings of the 44th Annual Design Automation Conference, DAC '07*, page 962–965, New York, NY, USA, 2007. Association for Computing Machinery.
  119. A. A. Melnikov et al. Active learning machine learns to create new quantum experiments. *Proceedings of the National Academy of Sciences*, 115(6):1221–1226, 2018.
  120. T. Menke, F. Häse, S. Gustavsson, A. J. Kerman, W. D. Oliver, and A. Aspuru-Guzik. Automated design of superconducting circuits and its application to 4-local couplers. *npj Quantum Information*, 7(1):49, 2021.
  121. B. Mete, M. Schulz, and M. Ruefenacht. Predicting the optimizability for workflow decisions. In *2022*

- IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)*, pages 68–74, 2022.
122. N. Meyer, C. Ufrecht, et al. A survey on quantum reinforcement learning, 2024.
  123. R. Miller. Multiverse computing raises 40 million to develop quantum software. *TechCrunch*, March 2024.
  124. N. Mohseni, T. Morstyn, C. O. Meara, D. Bucher, J. Nüßlein, and G. Cortiana. A competitive showcase of quantum versus classical algorithms in energy coalition formation, 2024.
  125. H. Moon et al. Machine learning enables completely automatic tuning of a quantum device faster than human experts. *Nature communications*, 11(1):4161, 2020.
  126. L. Moro, M. G. Paris, M. Restelli, and E. Prati. Quantum compiling by deep reinforcement learning. *Communications Physics*, 4(1):178, 2021.
  127. S. Naguleswaran, L. B. White, and I. Fuss. Automated planning using quantum computation. In *Proceedings of the Sixteenth International Conference on International Conference on Automated Planning and Scheduling*, ICAPS’06, page 418–421. AAAI Press, 2006.
  128. S. Nesbigall. Quantenbasierte koordination von multi-agentensystemen. Master’s thesis, Saarland University, Computer Science Department, 2008.
  129. N. M. P. Neumann et al. Multi-agent reinforcement learning using simulated quantum annealing. In *Computational Science – ICCS 2020*, pages 562–575, Cham, 2020. Springer International Publishing.
  130. M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
  131. S. Nolan, A. Smerzi, and L. Pezzè. A machine learning approach to bayesian parameter estimation. *npj Quantum Information*, 7(1):169, 2021.
  132. V. Oliveira Santos, F. P. Marinho, P. A. Costa Rocha, J. V. G. Thé, and B. Gharabaghi. Application of quantum neural network for solar irradiance forecasting: A case study using the folsom dataset, california. *Energies*, 17(14):3580, 2024.
  133. S. Ossowski and R. Menezes. On coordination and its significance to distributed and multi-agent systems. *Concurrency and Computation: Practice and Experience*, 18(4):359–370, 2006.
  134. A. Paler et al. Machine learning optimization of quantum circuit layouts. *ACM Transactions on Quantum Computing*, 4(2), feb 2023.
  135. I.-B. Park et al. A reinforcement learning approach to robust scheduling of semiconductor manufacturing facilities. *IEEE Transactions on Automation Science and Engineering*, 17(3):1420–1431, 2020.
  136. S. Park et al. Quantum multi-agent reinforcement learning for autonomous mobility cooperation. *IEEE Communications Magazine*, 2023.
  137. E. Pelofske, A. Bärtschi, and S. Eidenbenz. Quantum volume in practice: What users can expect from nisq devices. *IEEE Transactions on Quantum Engineering*, 3:1–19, 2022.
  138. B. Pokharel and D. A. Lidar. Better-than-classical grover search via quantum error detection and suppression. *npj Quantum Information*, 10(1):23, 2024.
  139. L. M. Possati. Ethics of quantum computing: An outline. *Philosophy & Technology*, 36(3):48, 2023.
  140. E. M. Pothos and J. R. Busemeyer. Quantum cognition. *Annual review of psychology*, 73(1):749–778, 2022.
  141. F. Pusse and M. Klusch. Hybrid online pomdp planning and deep reinforcement learning for safer self-driving cars. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1013–1020, 2019.
  142. A. Pyrkov et al. Quantum computing for near-term applications in generative chemistry and drug discovery. *Drug Discovery Today*, 28(8):103675, 2023.
  143. N. Quetschlich, L. Burgholzer, and R. Wille. Compiler optimization for quantum computing using reinforcement learning. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2023.
  144. N. Quetschlich, L. Burgholzer, and R. Wille. Predicting good quantum circuit compilation options. In *2023 IEEE International Conference on Quantum Software (QSW)*, pages 43–53, Los Alamitos, CA, USA, jul 2023. IEEE Computer Society.
  145. T. Rahwan. *Algorithms for coalition formation in multi-agent systems*. PhD thesis, University of Southampton, 2007.
  146. T. Rahwan, T. P. Michalak, M. Wooldridge, and N. R. Jennings. Coalition structure generation: A survey. *Artificial Intelligence*, 229:139–174, 2015.
  147. S. Rainjonneau et al. Quantum algorithms applied to satellite mission planning for earth observation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:7062–7075, 2023.
  148. R. Rake et al. Enterprise quantum computing market size, share, competitive landscape and trend analysis report, by component, deployment mode, technology and application, industry vertical: Global opportunity analysis and industry forecast, 2021-2030. [www.alliedmarketresearch.com/enterprise-quantum-computing-market](http://www.alliedmarketresearch.com/enterprise-quantum-computing-market), 2021.
  149. J. Rau. *Quantum theory: an information processing approach*. Oxford University Press, 2021.
  150. P. Rebentrost and S. Lloyd. Quantum computational finance: quantum algorithm for portfolio optimization. *KI-Künstliche Intelligenz*, pages 1–12, 2024.
  151. P. Rebentrost, M. Mohseni, and S. Lloyd. Quantum support vector machine for big data classification. *Phys. Rev. Lett.*, 113:130503, Sep 2014.
  152. J. Reichental. Quantum artificial intelligence is closer than you think. *Forbes*, November 20 2023.
  153. E. G. Rieffel and W. H. Polak. *Quantum computing: A gentle introduction*. MIT press, 2011.
  154. M. Rossignolo et al. Quocs: The quantum optimal control suite. *Computer Physics Communications*, 291:108782, 2023.
  155. A. S. Roy et al. Software tool-set for automated quantum system identification and device bring up, 2022.
  156. S. J. Russell and P. Norvig. *Artificial intelligence: a modern approach*. Pearson, 2016.
  157. S. Sarkar, M. Curado Malta, and A. Dutta. A survey on applications of coalition formation in multi-agent systems. *Concurrency and Computation: Practice and Experience*, 34(11):e6876, 2022.
  158. P. Schworm, X. Wu, M. Glatt, and J. C. Aurich. Solving flexible job shop scheduling problems in manufacturing with quantum annealing. *Production Engineering*, 17(1):105–115, 2023.
  159. M. B. Seelbach et al. Adiabatic quantum graph matching with permutation matrix constraints. In *2020 International Conference on 3D Vision (3DV)*, pages 583–592, Los Alamitos, CA, USA, nov 2020. IEEE Computer Society.
  160. P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

161. S. Sim, P. D. Johnson, and A. Aspuru-Guzik. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070, 2019.
162. A. Singh. What is quantum artificial intelligence? Posted on Medium.com on 17.7.2023, 2023.
163. A. Sinha, A. Macaluso, and M. Klusch. Nav-q: Quantum deep reinforcement learning for collision-free navigation of self-driving cars, 2023.
164. T. M. Stace et al. Optimized bayesian system identification in quantum devices. *Phys. Rev. Appl.*, 21:014012, Jan 2024.
165. P. Steinmüller, T. Schulz, F. Graf, and D. Herr. explainable ai for quantum machine learning, 2022.
166. M. P. V. Stenberg, O. Köhn, and F. K. Wilhelm. Characterization of decohering quantum systems: Machine learning approach. *Phys. Rev. A*, 93:012122, Jan 2016.
167. J. Sun and L. Xu. Cloud-based adaptive quantum genetic algorithm for solving flexible job shop scheduling problem. In *2019 IEEE 7th International Conference on Computer Science and Network Technology (ICCSNT)*, pages 1–5, 2019.
168. M. Swan et al. *Quantum Computing for the Brain*. WORLD SCIENTIFIC (EUROPE), 2021.
169. M. R. Team. Quantum computing in automotive market by application deployment, component, stakeholder & region - global forecast to 2035. *GII Research*, 2023.
170. F. Tennie and T. N. Palmer. Quantum computers for weather and climate prediction: The good, the bad, and the noisy. *Bulletin of the American Meteorological Society*, 104(2):E488–E500, 2023.
171. D. Thanos et al. Automated reasoning in quantum circuit compilation. In *Model Checking Software (SPIN)*, 2024.
172. E. B. Unlu, M. Comajoan Cara, G. R. Dahale, Z. Dong, R. T. Forestano, S. Gleyzer, D. Justice, K. Kong, T. Magorsch, K. T. Matchev, et al. Hybrid quantum vision transformers for event classification in high energy physics. *Axioms*, 13(3):187, 2024.
173. S. E. Venegas-Andraca, S. Bose, et al. Quantum computation and image processing: New trends in artificial intelligence. In *IJCAI*, volume 1563, 2003.
174. S. M. Venkatesh et al. Q-seg: Quantum annealing-based unsupervised image segmentation. *IEEE Computer Graphics and Applications*, pages 1–13, 2024.
175. S. M. Venkatesh, A. Macaluso, and M. Klusch. Bilp-q: quantum coalition structure generation. In *Proceedings of the 19th ACM International Conference on Computing Frontiers*, CF '22, page 189–192, New York, NY, USA, 2022. Association for Computing Machinery.
176. S. M. Venkatesh, A. Macaluso, and M. Klusch. Gcs-q: Quantum graph coalition structure generation. In *Computational Science – ICCS 2023: 23rd International Conference, Prague, Czech Republic, July 3–5, 2023, Proceedings, Part V*, page 138–152, Berlin, Heidelberg, 2023. Springer-Verlag.
177. S. M. Venkatesh, A. Macaluso, and M. Klusch. Quacs: Variational quantum algorithm for coalition structure generation in induced subgraph games. In *Proceedings of the 20th ACM International Conference on Computing Frontiers*, CF '23, page 197–200, New York, NY, USA, 2023. Association for Computing Machinery.
178. S. M. Venkatesh, A. Macaluso, M. Nuske, M. Klusch, and A. Dengel. Quantum annealing-based algorithm for efficient coalition formation among leo satellites. *arXiv preprint arXiv:2408.06007*, 2024.
179. S. M. Venkatesh, A. Macaluso, M. Nuske, M. Klusch, and A. Dengel. Qubit-efficient variational quantum algorithms for image segmentation. *arXiv preprint arXiv:2405.14405*, 2024.
180. J. Wallnöfer et al. Machine learning for long-distance quantum communication. *PRX Quantum*, 1:010301, Sep 2020.
181. Y. Wang, X. Ma, Z. Li, Y. Liu, M. Xu, and Y. Wang. Profit distribution in collaborative multiple centers vehicle routing problem. *Journal of Cleaner Production*, 144:203–219, 2017.
182. S. Warwas et al. Bochica: A model-driven framework for engineering multiagent systems. In *ICAART (1)*, pages 109–118, 2012.
183. M. Weigold et al. Encoding patterns for quantum algorithms. *IET Quantum Communication*, 2(4):141–152, 2021.
184. G. Weiss. *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT press, 1999.
185. D. Widdows et al. Quantum natural language processing, 2024.
186. N. Wittler et al. Integrated tool set for control, calibration, and characterization of quantum devices applied to superconducting qubits. *Phys. Rev. Appl.*, 15:034080, Mar 2021.
187. A. Wozniakowski et al. Boosting on the shoulders of giants in quantum device calibration, 2020.
188. L.-L. Ye, C. Arenz, J. M. Lukens, and Y.-C. Lai. Entanglement engineering of optomechanical systems by reinforcement learning, 2024.
189. W. J. Yun et al. Quantum multiagent actor-critic neural networks for internet-connected multirobot coordination in smart factory management. *IEEE Internet of Things Journal*, 10(11):9942–9952, 2023.
190. W. Zeng and B. Coecke. Quantum algorithms for compositional natural language processing. *arXiv preprint arXiv:1608.01406*, 2016.
191. J. Zhang et al. Review of job shop scheduling research and its new perspectives under industry 4.0. *Journal of intelligent manufacturing*, 30:1809–1830, 2019.
192. Q. Zhang and S. Hu. An improved hybrid quantum particle swarm optimization algorithm for fjsp. In *Proceedings of the 2019 11th International Conference on Machine Learning and Computing*, ICMLC '19, page 246–252, New York, NY, USA, 2019. Association for Computing Machinery.
193. Y. Zhuang et al. Quantum computing in intelligent transportation systems: A survey, 2024.
194. S. Zibaei, A. Hafezalkotob, and S. S. Ghashami. Cooperative vehicle routing problem: an opportunity for cost saving. *Journal of Industrial Engineering International*, 12:271–286, 2016.