
Variational Algorithms in Quantum Artificial Intelligence:

Training Optimization, Automated Circuit
Design and Applications

Michael Kölle



2025

Variational Algorithms in Quantum Artificial Intelligence:

Training Optimization, Automated Circuit
Design and Applications

Michael Kölle

Dissertation
an der Fakultät für Mathematik, Informatik und Statistik
der Ludwig-Maximilians-Universität München

vorgelegt von
Michael Kölle
aus Dillingen an der Donau

München, 10. Dezember 2025

1. Gutachter:	Prof. Dr. Claudia Linnhoff-Popien
2. Gutachter:	Prof. Dr. Stefan Fischer
3. Gutachter:	Prof. Dr. Jörg Hähner
Tag der mündlichen Prüfung:	15. Juli 2026

Eidesstattliche Versicherung

Hiermit erkläre ich an Eides statt, dass die vorliegende Arbeit selbständig und ohne fremde Hilfe verfasst und keine anderen als die angegebenen Hilfsmittel verwendet wurden. Die Stellen der Arbeit, die dem Wortlaut oder dem Sinn nach anderen Werken entnommen sind, wurden unter Angabe der Quellen kenntlich gemacht. Die Kennzeichnung gilt auch für diejenigen Stellen, die mit Hilfe von KI-Werkzeugen erstellt oder bearbeitet wurden. Sie gehen in vollem Umfang auf mein geistiges Eigentum zurück und verletzen nicht das geistige Eigentum Dritter.

München, 10. Dezember 2025

Michael Kölle

This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>

Acknowledgments

This thesis marks the conclusion of a long and exciting journey, which was only possible with the support, guidance, and encouragement of many people. I would like to thank all those who have accompanied me on this path.

My deepest gratitude goes to my supervisor, Prof. Claudia Linnhoff-Popien. I thank her for her unwavering trust and the freedom she granted me in my research. I am particularly grateful for the chance to lead the technical department and our AIONA lab, to oversee the production of our yearly DIGICON event, and to establish the “Quantum Artificial Intelligence & Optimization” workshop. The experiences I gained in teaching—especially in our award-winning lectures “Computer Architecture” and “Operating Systems”, as well as countless theses, seminars and practicals—were invaluable to me. I am equally grateful for the opportunity to contribute to PlanQK, QAR-Lab Bayern, Innomi, CAQAO, Munich Quantum Valley, and QuCUN, gaining insight into quantum computing research across academia and industry.

I extend my heartfelt thanks to my examination committee for their time, dedication and constructive review: the second examiner, Prof. Stefan Fischer; the third examiner, Prof. Jörg Hähner; the chairperson, Prof. Matthias Schubert; and the substitute examiner, Prof. Caroline Friedel.

A PhD is seldom the work of an individual, but the result of a collaborative environment. I sincerely thank the colleagues who warmly welcomed me: Steffen Illium, Robert Müller, Kyrill Schmid, and Fabian Ritz. Special thanks to Thomas Gabor and Thomy Phan for being scientific role models and mentors. It was a pleasure to co-teach the aforementioned lectures with Julian Hager, the “Quantum Applications” and “Intelligent Systems” lectures with Jonas Nüßlein and the QCP practical course with Jonas Stein and Max Zorn. The conference trips with Philipp Altmann and Leo Sünkel will remain unforgettable. Finally, I thank Gerhard Stenzel, Tobias Rohe, Sebastian Wölckert, Alexander Klingebiel, Nina Freise, Alexander Feist, and Simon Schlichting for our shared joy and success in “AI for Health” research and teaching. The chair would not function without its organizational core. I thank Melisa Delic, Katja Grenner, and Magdalena Gendritzki for their tireless dedication. A big thank you to our technical team—Tim Bücher, Paul Specht, Aleksander Szymanski and David Galesic—for their fantastic support.

Most importantly, my gratitude goes to my family and friends. Your love, patience, and unconditional support were the foundation upon which this success is built. Thank you for always being there for me.

Zusammenfassung

In den letzten Jahren wurde der Fortschritt der Künstlichen Intelligenz (KI) maßgeblich durch die Skalierung der Modellgröße und -komplexität vorangetrieben. Dieser stetige Anstieg der Parameterzahlen und der daraus resultierende enorme Ressourcenverbrauch stellen die Skalierbarkeit klassischer Architekturen infrage und erfordern fundamental neue, parametereffiziente Ansätze. Das Quantencomputing stellt hierbei eine vielversprechende Alternative dar, die leistungsfähige Modelle mit signifikant weniger Parametern ermöglicht. In der heutigen Ära des Noisy Intermediate-Scale Quantum Computing etablieren sich variationelle Quantenalgorithmen (VQAs) als führender Ansatz für reale Hardware. Ihre Anwendung wird jedoch durch mehrere grundlegende Herausforderungen eingeschränkt: instabile Trainingsdynamiken, flache Lösungslandschaften, welche die Skalierbarkeit erschweren, sowie den mühsamen manuellen Entwurf von Quantenschaltungen.

Diese Dissertation begegnet diesen Herausforderungen mit der Entwicklung systematischer Methodiken für den gesamten VQA-Entwicklungszyklus. Es werden robuste Trainings- und Optimierungstechniken wie Weight Re-Mapping, exponentielle Lernratenanpassung, strukturelles Pruning und metaheuristische Optimierung eingeführt, um Konvergenz zu steigern und flache Lösungslandschaften zu vermeiden. Zudem werden Verfahren zur automatisierten Architektursuche für Quantenschaltungen mittels Reinforcement Learning und evolutionärer Algorithmen entwickelt, um die Abhängigkeit von manuellen Entwürfen zu verringern. Der praktische Nutzen dieser Methodik wird in mehreren anspruchsvollen KI-Domänen demonstriert. So werden VQAs genutzt, um hochgradig parametereffiziente Agenten für das Quanten-Reinforcement-Learning zu konstruieren. Ferner werden neuartige Quantenmodelle für generative KI und die Anomalieerkennung vorgestellt. Über diese Anwendungen hinaus wird eine Methode zur fairen Bewertung hybrider Systeme entwickelt und das Verhältnis zwischen Parametereffizienz und Trainingszeit quantifiziert. Zusammengefasst ergeben diese Beiträge validierte, systematische Ansätze, die VQAs von theoretischen Konzepten zu robusten, praxistauglichen Werkzeugen für den Einsatz von Quanten-KI auf Hardware der nahen Zukunft weiterentwickeln.

Abstract

In recent years, the progress of artificial intelligence (AI) has been largely driven by scaling the size and complexity of its models. This strategy, however, has led to ever-increasing parameter counts and unsustainable resource demands, raising questions about the future scalability of classical architectures and creating a clear need for more parameter-efficient approaches. Quantum computing offers a promising alternative, capable of creating powerful models with significantly fewer parameters. In the current noisy intermediate-scale quantum (NISQ) era, variational quantum algorithms (VQAs) have emerged as the leading strategy to realize this potential on near-term hardware. However, several foundational challenges limit the practical use of VQAs: unstable training dynamics, the barren plateau phenomenon, which inhibits scalability, and the manual, heuristic-based process of circuit design.

In this thesis, we address these challenges by developing a systematic framework for the entire VQA development lifecycle. We introduce robust training and optimization techniques—including weight re-mapping, exponential learning rate decay, structural pruning, and metaheuristic optimization—to ensure stable convergence and mitigate barren plateaus. Furthermore, we establish methods for automated circuit architecture search using reinforcement learning and evolutionary algorithms to reduce the reliance on manual design. We demonstrate the practical value of this framework in several demanding AI domains. We use VQAs to create highly parameter-efficient quantum reinforcement learning agents and introduce novel quantum models for generative AI and anomaly detection. Beyond these applications, we provide an analysis of the models’ real-world trade-offs, establishing a method to fairly assess hybrid systems and quantifying the balance between parameter efficiency and training time. Taken together, these contributions provide a set of validated, systematic approaches that advance VQAs from theoretical concepts toward robust, practical tools for near-term quantum AI.

Contents

1. Introduction	1
1.1. Motivation	1
1.2. Scope and Research Questions	2
1.3. Thesis Structure	4
2. Foundations	7
2.1. Artificial Intelligence	7
2.1.1. Supervised Learning	8
2.1.2. Reinforcement Learning	9
2.1.3. Evolutionary and Heuristic Optimization	11
2.2. Quantum Computing	12
2.2.1. Quantum-Mechanical Principles	12
2.2.2. Quantum Gates	14
2.2.3. Quantum Circuits	15
2.3. Quantum Artificial Intelligence	16
2.3.1. Variational Quantum Algorithms	17
2.3.2. Variational Quantum Circuits	18
3. Training Optimization Techniques	21
3.1. Adapting Classical Optimization Paradigms	23
3.2. Training Heuristics	24
3.3. Strategies for Mitigating Barren Plateaus	25
4. Automated Circuit Design and Synthesis	29
4.1. State Preparation using Reinforcement Learning	30
4.2. Ansatz Design using Reinforcement Learning	31
4.3. Ansatz Design using Evolutionary Strategies	33
4.4. Mutation Strategies for Circuit Design	35
5. Quantum Reinforcement Learning	37
5.1. Hybrid Quantum-Classical Actor-Critic Agents	38
5.2. Metaheuristic Optimization for QRL Agents	39
5.3. Evolutionary Multi-Agent Quantum Reinforcement Learning	40
6. Advanced Models and Applications	43
6.1. Quantum Denoising Diffusion Models	44
6.2. Efficient Quantum Anomaly Detection	46

6.3.	Disentangling Hybrid Quantum-Classical Models	48
6.4.	Comparative Analysis of Parameter Efficiency	50
7.	Conclusion	53
7.1.	Summary	53
7.2.	Outlook	56
	Acronyms	57
	List of Figures	61
	Bibliography	63
	Appendix	83
A.	Publications	83
A.0.	Overview	84
A.1.	Towards Efficient Quantum Anomaly Detection: One-Class SVMs Using Variable Subsampling and Randomized Measurements	88
A.2.	Evaluating Mutation Techniques in Genetic Algorithm-Based Quantum Circuit Synthesis	89
A.3.	Evaluating Parameter-Based Training Performance of Neural Net- works and Variational Quantum Circuits	90
A.4.	Improving Convergence for Quantum Variational Classifiers Using Weight Re-Mapping	91
A.5.	Weight Re-mapping for Variational Quantum Algorithms	92
A.6.	Quantum Advantage Actor-Critic for Reinforcement Learning . . .	93
A.7.	Investigating the Lottery Ticket Hypothesis for Variational Quan- tum Circuits	94
A.8.	Quantum Transfer Learning: The Impact of Classical Preprocessing	95
A.9.	Disentangling Quantum and Classical Contributions in Hybrid Quan- tum Machine Learning Architectures	96
A.10.	Quantum Architecture Search for Solving Quantum Machine Learn- ing Tasks	97
A.11.	Architectural Influence on Variational Quantum Circuits in Multi- Agent Reinforcement Learning: Evolutionary Strategies for Opti- mization	98
A.12.	A Reinforcement Learning Environment for Directed Quantum Cir- cuit Synthesis	99
A.13.	Optimizing Variational Quantum Circuits Using Metaheuristic Strate- gies in Reinforcement Learning	100
A.14.	Quantum Denoising Diffusion Models	101
A.15.	Multi-Agent Quantum Reinforcement Learning Using Evolution- ary Optimization	102

A.16. A Study on Optimization Techniques for Variational Quantum Circuits in Reinforcement Learning	103
---	-----

1. Introduction

We cannot solve our problems with the same thinking we used when we created them.

Albert Einstein (attributed [149])

1.1. Motivation

Artificial Intelligence (AI), particularly Deep Learning (DL), has achieved major successes in the last decade, delivering state-of-the-art performance in domains ranging from scientific discovery [65] and medical diagnostics [41] to complex decision-making [134] and natural language processing [19, 152]. These breakthroughs have led to widespread use and economic growth [7], making AI a key part of modern life. Much of this progress has been driven by a strategy of scaling, achieving performance gains by dramatically increasing model complexity. The number of trainable parameters, especially in transformer-based generative models, has expanded from millions to hundreds of billions, establishing a clear correlation between a model’s scale and its capabilities [19, 67]. However, this reliance on scaling has raised concerns about its long-term viability. The computational resources required to train these massive models have grown faster than improvements in hardware efficiency, leading to questions about the future scalability of classical architectures [5, 129]. While the economic benefits currently justify the investment, the escalating resource demands suggest that continued progress cannot rely on scaling alone and call for new, more parameter-efficient computational paradigms.

Quantum Computing (QC) presents such an inherently different approach to information processing, rooted in the principles of quantum mechanics [44]. By using phenomena such as superposition and entanglement, quantum computers are capable of navigating and manipulating information within an exponentially large computational state space [110]. A register of just n quantum bits (qubits) can represent 2^n classical states simultaneously, offering a potential pathway to algorithmic speedups for certain classes of problems that are intractable for even the most powerful classical supercomputers [110]. For AI, this paradigm is particularly promising, as quantum algorithms have the potential to improve Machine Learning (ML) by executing complex linear algebra

operations in high-dimensional Hilbert spaces, thereby offering a potential solution to the scaling and resource bottlenecks that currently constrain classical DL [15]. Despite its theoretical promise, the practical realization of large-scale, fault-tolerant QC remains a distant goal. The current landscape is defined by the Noisy Intermediate-Scale Quantum (NISQ) era, characterized by hardware with a limited number of qubits, high sensitivity to environmental noise and a lack of robust quantum error correction [118]. In this constrained reality, a purely quantum-based approach is infeasible for complex AI tasks. Consequently, a pragmatic compromise has emerged in the form of Variational Quantum Algorithms (VQAs) [23]. VQAs employ a hybrid architecture, leveraging shallow¹, parameterized Variational Quantum Circuits (VQCs) (similar to classical Artificial Neural Networks (NNs)) to perform computationally intensive tasks while relying on classical optimizers to navigate the noisy hardware and train the model’s parameters. Although it does not have a proven computational speedup, this hybrid approach represents the most viable and promising strategy for using the power of near-term quantum devices for practical Quantum Machine Learning (QML) applications.

Though VQAs present a promising direction, their transition from theoretical concepts to robust, practical tools is hindered by a set of unaddressed challenges that can be grouped into three core areas. First, at the level of training and optimization, their development is notoriously difficult [99]. The training process is often unstable, a difficulty increased by the barren plateau phenomenon, where gradients vanish exponentially with system size, rendering gradient-based training ineffective for all but the smallest models [99]. Second, in architectural design, the structure of a VQC—its ansatz—is a primary factor influencing its performance, yet the design of these circuits remains a largely manual and heuristic process, creating a significant bottleneck in discovering optimal, task-specific structures [66]. Finally, the field lacks methodological consistency in its evaluation [17]. In many hybrid systems, it is unclear how much each component contributes to performance, making it difficult to assess claims of quantum superiority. Furthermore, there is a distinct lack of systematic studies that quantify the real-world trade-off between the claimed parameter efficiency of VQCs and their training costs.

1.2. Scope and Research Questions

In this thesis, we systematically confront the aforementioned challenges in training, design and evaluation, aiming to mature VQAs from a promising theoretical framework into a set of practical and effective tools for ML. To this end, we present a series of contributions that address the full VQA development lifecycle. We begin by developing techniques to improve the stability

¹In this context, shallow refers to circuits with a low depth, meaning a small number of sequential gate operations.

and robustness of the training process, and then establish methods for automated circuit synthesis to replace manual design. These enhanced methods are then applied to demanding domains like Quantum Reinforcement Learning (QRL) and generative modeling to test their efficacy. Finally, we provide a systematic, data-driven analysis of the performance and parameter efficiency of VQAs to build a more methodologically sound approach to Quantum Artificial Intelligence (QAI). This research agenda is guided by the following six questions:

- Q1 How can VQAs in the NISQ era be designed to mitigate data capacity and convergence limitations?** This question is motivated by the need to make VQAs practical on current hardware. We investigate hybrid architectures that use classical dimensionality reduction to manage the limited qubit count of current quantum hardware, enabling the processing of high-dimensional data. We also explore methods to stabilize the often unreliable training process by adapting classical optimization techniques to the unique properties of quantum circuits and by systematically evaluating training heuristics.
- Q2 How can VQAs be made more robust and resilient to challenges like barren plateaus?** This question examines strategies to overcome the barren plateau phenomenon, a major obstacle to scaling VQAs. In this work, we investigate two different approaches: first, circumventing the vanishing gradient problem by employing gradient-free metaheuristic algorithms for optimization and second, simplifying the model’s complexity by applying structural pruning to remove redundant parameters and identify smaller, more trainable subcircuits.
- Q3 How can the manual process of VQC design be automated to discover more efficient, task-specific architectures?** This question is motivated by the limitations of manual, intuition-driven quantum circuit design, which is a labor-intensive process constrained by human heuristics. It explores the development of automated methods that can systematically search the large space of possible circuit architectures. In this work, we examine two ML approaches for this task: Reinforcement Learning (RL) and Evolutionary Algorithms (EAs). The goal is to discover novel and efficient circuit designs.
- Q4 How can VQCs be effectively integrated into QRL agents to address challenges in stability, optimization and scalability?** This question explores applying QC to the computationally demanding domain of RL. The motivation is to determine if the properties of VQCs can enhance the performance of RL agents. We examine several challenges by developing hybrid quantum-classical agents to improve training stability, evaluating metaheuristic algorithms for optimization and using the parameter efficiency of VQCs to show a path toward scalable quantum multi-agent systems.

- Q5 How can quantum models be designed and adapted to address advanced generative and discriminative ML tasks?** This question is motivated by the goal of testing quantum models on advanced ML tasks beyond simple classification and RL. For generative modeling, we explore the design of parameter-efficient quantum circuits for diffusion models, a leading class of generative architectures. For discriminative tasks, we examine the computational scaling of quantum kernel methods in anomaly detection, developing techniques to make them more suitable for real-world datasets.
- Q6 What is the empirical trade-off between the parameter efficiency of VQCs and their training time compared to classical NNs?** This question aims to quantitatively evaluate a central claim of QML: that quantum models can achieve comparable performance with fewer parameters. The motivation is to move beyond theoretical arguments and provide a data-driven analysis of the practical advantages and current limitations of VQCs. This is done through a systematic comparison where we analyze the differences in parameter counts and training durations between performance-matched VQCs and classical NNs.

These six questions define a comprehensive research agenda that addresses the end-to-end VQA development process from foundational training mechanics and automated design to high-level application performance and empirical analysis. They are not addressed in isolation but form an integrated study into building more effective QML models. The remainder of this thesis is structured to systematically answer these questions, with each chapter building upon the last to construct a complete and practical framework for QAI.

1.3. Thesis Structure

We begin by establishing the necessary theoretical groundwork in AI and QC in chapter 2. With this foundation in place, chapter 3 directly confronts important VQA training challenges, introducing techniques to improve stability and presenting methods to mitigate barren plateaus. Building on this, chapter 4 moves beyond manual heuristics by detailing automated circuit synthesis frameworks using RL and EAs.

We then demonstrate the practical efficacy of these enhanced methods. chapter 5 explores the demanding domain of QRL, highlighting the parameter efficiency of VQC-based agents and validating gradient-free optimizers. Chapter 6 broadens this scope to more advanced tasks, including generative diffusion modeling and scalable anomaly detection. This chapter also introduces an architecture for better differentiating the contributions of quantum and classical components in hybrid models. Finally, we summarize the key findings, synthesize the contributions and discuss future directions for QAI in chapter 7.

Table 1.1.: An overview of the main technical contents of this thesis, research questions and publications.

Thesis chapter	Question	Publications
3. Training Optimization Techniques	Q1	75, 76, 78–80, 87
	Q2	75, 76, 78, 84
4. Automated Circuit Design and Synthesis	Q3	73, 81–83
5. Quantum Reinforcement Learning	Q4	77, 84, 86, 87
6. Advanced Models and Applications	Q5	72, 79, 80, 85
	Q6	74

Table 1.1 maps the main chapters of this thesis to the specific research questions they address. To maintain a focused narrative, the following chapters summarize the key findings from each publication. The full text of each paper, which includes a more comprehensive review of related work as well as detailed information on credits, author contributions, and the purpose of each study, is provided in appendix A.

2. Foundations

Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical.

Richard P. Feynman, 1981 [44]

2.1. Artificial Intelligence

AI has pursued the scientific goal of equipping machines with the ability to think, learn and act autonomously—qualities that were long considered exclusive to human cognition. However, the idea predates modern computers: in his seminal essay on AI in 1950, Alan Turing posed the famous question “Can machines think?” [151]. Six years later, in 1956, AI was declared an official scientific discipline in the Dartmouth Workshop organized by McCarthy, Minsky, Shannon and others. Initial optimism focused on symbolic reasoning and rule-based systems [20, 107]. However, the field experienced alternating phases of rapid progress and so-called AI winters [31, 121], which were mainly caused by inappropriate expectations and computational limitations. The late 1980s saw a resurgence of the field, rekindling interest in adaptive systems capable of learning from data—in particular, the backpropagation algorithm for training multilayer perceptrons [120]. The following decades brought exponential growth in data availability, computing power [5, 104] and algorithm design [92].

These conditions favored ML as a data-driven approach that enables machines to improve through experience rather than explicit programming. Any machine learning paradigm—be it supervised, unsupervised, or reinforcement—can be formally described by four elements. It consists of a *dataset* $\mathcal{D} = \{z_i\}_{i=1}^N \subseteq \mathcal{Z}$, which is drawn independently from an unknown distribution $P_{\mathcal{Z}}$. A *model*, a parametrized hypothesis family $\{f_{\theta}\}_{\theta \in \Theta}$, with θ denoting the set of trainable parameters—typically, in modern ML, a NN acting as a universal function approximator¹ [32, 60]; or in QAI, this role can also be

¹A universal function approximator is a model that, given enough capacity (e.g., neurons), can approximate any continuous function to an arbitrary degree of accuracy.

fulfilled by a VQC as described in section 2.3.2. A *loss function* ℓ is chosen to quantify the difference between model predictions and desirable outcomes. An *optimization procedure* selects parameters that minimize the expected risk $\mathcal{R}(f_\theta) = \mathbb{E}_{z \sim P_Z} [\ell(f_\theta(z))]$. Because the underlying distribution P_Z is by definition unknown, the expected risk is approximated by the empirical risk $\hat{\mathcal{R}}_N(f_\theta) = \frac{1}{N} \sum_{i=1}^N \ell(f_\theta(z_i))$, defined over the dataset $\mathcal{D}$. In the supervised setting a sample of $\mathcal{D}$ is of the form $z = (x, y)$, whereas in an unsupervised setting the label y is absent; RL, by contrast, operates on experience tuples (s, a, r, s') as discussed in the following sections.

A particularly successful field of ML is DL, where deep NNs with many layers of nonlinear transformations learn hierarchical representations directly from raw data. Concretely, an L -layer feed-forward network implements the mapping $f_\theta(x) = g_L \circ g_{L-1} \circ \dots \circ g_1(x)$, $g_\ell(z) = \sigma_\ell(W_\ell z + b_\ell)$, where $x \in \mathcal{X} \subseteq \mathbb{R}^{d_0}$ is the input, $W_\ell \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$ is the weight, $b_\ell \in \mathbb{R}^{d_\ell}$ is the corresponding bias and $\theta = \{W_\ell, b_\ell\}_{\ell=1}^L$ denotes the trainable parameters. A non-linear activation σ_ℓ (e.g., Rectified Linear Unit (ReLU) [106], Gaussian Error Linear Unit (GELU) [58], or sigmoid [153]) makes the overall function non-linear in its input and parameters. The depth of the network allows it to express hierarchical features: early layers capture (local) low-level patterns, whereas deeper layers learn increasingly abstract representations. Under mild conditions, even a single hidden layer network is a universal approximator [32, 60]. Nevertheless, stacking many layers often leads to a dramatic improvement of statistical and computational efficiency [101]. Training deep NNs requires stochastic optimization in a high-dimensional, non-convex landscape. Let $\mathcal{B}_t \subset \mathcal{D}$ be a mini-batch of cardinality B . Stochastic Gradient Descent (SGD) updates the parameters at iteration t using the gradient of $\hat{\mathcal{R}}_N(f_\theta)$ via the back-propagation algorithm [120] according to $\theta_{t+1} = \theta_t - \eta_t \nabla_\theta [\frac{1}{B} \sum_{(x,y) \in \mathcal{B}_t} \ell(f_\theta(x))]$, where $\eta_t > 0$ denotes the learning rate. The vanilla SGD can be improved through momentum [142] or adaptive learning-rate schedules, such as Adaptive Moment Estimation (Adam) [69] or Root Mean Square Propagation (RMSProp) [150].

2.1.1. Supervised Learning

In Supervised Learning (SL), a model learns a mapping from inputs to desired outputs using a set of labeled examples. Typical tasks include classification (predicting a discrete label) and regression (predicting a continuous value). Formally, the problem is defined as follows. Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N \subseteq \mathcal{X} \times \mathcal{Y}$ be an independent and identically distributed sample from an unknown distribution $P_{\mathcal{X} \times \mathcal{Y}}$. The goal is to find parameters θ of a hypothesis $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ that minimize the expected risk $\mathcal{R}(f_\theta)$; in practice we optimize its empirical counterpart $\hat{\mathcal{R}}_N(f_\theta)$. A model is then trained and its parameters θ are updated at iteration t using a mini-batch $\mathcal{B}_t \subset \mathcal{D}$ according to: $\theta_{t+1} = \theta_t - \eta_t \nabla_\theta \hat{\mathcal{R}}_{|\mathcal{B}_t|}(f_\theta)|_{\theta=\theta_t}$, where η_t is the learning rate. For model evaluation,

domain-specific metrics are used. For classification, accuracy, precision, recall, F_1 -score and area under the Receiver Operating Characteristic (ROC) curve (Area Under the Curve (AUC)) are common [116]. Regression tasks often use Mean Absolute Error (MAE), Mean Squared Error (MSE) or Root Mean Squared Error (RMSE) [24].

2.1.2. Reinforcement Learning

Modern RL emerged in the 1980s from three major research fields: optimal control [11], trial-and-error learning [145] and Temporal-Difference (TD) learning [144]. It aims to solve sequential-decision-making problems under uncertainty. In RL, an agent learns to make decisions by interacting with an environment [143], receiving feedback in the form of rewards and improving its behavior over time in order to maximize its cumulative reward. Markov Decision Processes (MDPs) [62] provide a standardized framework for formalizing sequential decision-making scenarios, representing the interaction between an agent and its environment as states, actions, transitions and rewards.

Markov Decision Processes A sequential decision-making scenario can be formalized as an MDP [62]. An MDP is defined as a 5-tuple $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{S}$ is the set of possible states, $\mathcal{A}$ is the set of available actions, $P(s' | s, a)$ is the state transition probability function specifying the likelihood of reaching state s' from state s after taking action a , $R(s, a, s')$ is the reward function that assigns a scalar feedback for each transition and $\gamma \in [0, 1]$ is the discount factor that weighs future rewards. Within this framework, the agent follows a policy $\pi(a | s)$, which defines the probability of taking action a in state s . The goal of the agent is to learn a policy $\pi(a | s)$ that maximizes the expected cumulative reward $\mathbb{E}[G_t] = \mathbb{E}[\sum_{i=0}^{\infty} \gamma^i \cdot r_{t+i}]$. The state-value function $V^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a | s) \sum_{s' \in \mathcal{S}} P(s' | s, a) [R(s, a, s') + \gamma V^\pi(s')]$ evaluates the expected return starting from state s and following policy π thereafter. The action-value function $Q^\pi(s, a) = \sum_{s' \in \mathcal{S}} P(s' | s, a) [R(s, a, s') + \gamma \sum_{a' \in \mathcal{A}} \pi(a' | s') Q^\pi(s', a')]$ represents the expected return of taking action a in state s and subsequently following policy π . An important objective in RL is to find an optimal policy π^* that maximizes the expected return from every state. The optimal state-value function $V^*(s) = \max_{a \in \mathcal{A}} \sum_{s' \in \mathcal{S}} P(s' | s, a) [R(s, a, s') + \gamma V^*(s')]$ and action-value function $Q^*(s, a) = \sum_{s' \in \mathcal{S}} P(s' | s, a) [R(s, a, s') + \gamma \max_{a' \in \mathcal{A}} Q^*(s', a')]$ satisfy the Bellman optimality equations.

Temporal-Difference Learning While exact solutions to the Bellman equations require knowledge of the transition and reward functions, TD learning methods estimate value functions directly from sampled experience. Value estimates are updated based on the difference between successive predictions using the TD error $\delta_t = r_{t+1} + \gamma V(s_{t+1}) - V(s_t)$ in the TD(0) update rule

$V(s_t) \leftarrow V(s_t) + \alpha \delta_t$ [144], where α is the learning rate. TD methods enable incremental learning techniques that form the basis for many RL algorithms, including value-based methods like Q-Learning (QL), which is discussed in the next section.

Value-Based Methods Value-based approaches attempt to approximate the optimal action-value function $Q^*(s, a)$, from which the optimal policy $\pi^*(s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a)$ can be derived by selecting the action with the highest value in each state. These methods do not require an explicit model of the environment and instead update estimates from experience. One common value-based algorithm is *QL* [154]. QL employs a greedy action selection strategy, choosing the action that maximizes the estimated future return. The update rule is given by $Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r_{t+1} + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$. This forms the basis for extensions such as *Deep Q-Networks (DQNs)* [103], which approximate Q -values using NNs or VQCs in high-dimensional or continuous state spaces.

Policy Gradient Methods In contrast to value-based methods that learn a value function and derive a policy from it, policy gradient methods directly parametrize and optimize the policy itself. The policy is represented by a function $\pi_\theta(a | s)$ with parameters θ . The core idea is to perform stochastic gradient ascent on an objective function $J(\theta) = \mathbb{E}_{\pi_\theta}[G_t]$ that measures the expected cumulative reward of the policy π_θ [146]. Its gradient $\nabla_\theta J(\theta) = \mathbb{E}_{\pi_\theta}[\nabla_\theta \log \pi_\theta(a_t | s_t) Q^{\pi_\theta}(s_t, a_t)]$, where $\nabla_\theta \log \pi_\theta(a_t | s_t)$ is the score function and $Q^{\pi_\theta}(s_t, a_t)$ is the action-value function, is used to adjust the parameters θ of the policy to increase the likelihood of actions that lead to higher returns. To improve the stability and efficiency of this optimization, the action-value function $Q^{\pi_\theta}(s_t, a_t)$ is often replaced by the advantage function [102, 127], $A(s_t, a_t) = Q(s_t, a_t) - V(s_t)$, which quantifies how much better a particular action is compared to the average action at a given state. This substitution reduces variance in the gradient estimates and typically results in more stable and faster learning.

Actor-Critic Methods Actor-Critic methods unite the strengths of policy-based and value-based approaches by maintaining two separate structures: an *actor*, which updates the policy $\pi_\theta(a | s)$ in the direction suggested by policy gradients and a *critic*, which estimates the value function $V^\pi(s)$ or action-value function $Q^\pi(s, a)$ to provide feedback on the actor's performance. The critic helps reduce the variance of policy gradient estimates by serving as a baseline, while the actor uses this information to improve the policy. Typically, the advantage function $A(s_t, a_t)$ is computed using the critic's estimates and is used to weight the policy gradient update. This structure allows for more sample-efficient and stable learning, especially in environments with

high-dimensional or continuous action spaces. Popular implementations include *Advantage Actor-Critic (A2C)* [102], which uses synchronous updates across multiple agents to stabilize training and *Proximal Policy Optimization (PPO)* [128], which introduces a clipped objective to keep policy updates within a safe limit.

2.1.3. Evolutionary and Heuristic Optimization

While gradient-based methods such as SGD form the basis of modern ML, their effectiveness depends on the existence of a well-defined, differentiable objective function. Evolutionary and heuristic optimization methods offer a powerful class of gradient-free alternatives that work without these requirements. This family of methods includes classic single-solution heuristics such as Simulated Annealing (SA) [70] and Tabu Search (TS) [49], as well as population-based EAs such as Genetic Algorithms (GAs). Their usefulness is most notable for problems with non-differentiable or discontinuous loss functions, discrete or mixed search spaces and complex optimization landscapes. Such landscapes, which are full of local minima or phenomena such as barren plateaus, pose a well-known challenge in areas such as QAI [99]. Therefore, these methods are frequently used in ML for tasks such as architecture search, hyperparameter tuning and the direct optimization of VQCs, where gradient-based approaches may be infeasible or perform poorly. Most heuristics and EAs can be explained by a set of core functional components that define their search behavior.

Solution Representation The first component is the solution representation. Here, a candidate solution is encoded as a point x within a defined search space $\mathcal{X}$ and the structure of $\mathcal{X}$ usually depends on the problem. Examples include binary vectors $x \in \{0,1\}^n$ for subset selection or permutations for ordering problems [50]. In a GA, this representation is often referred to as the “chromosome” of an individual.

Objective Function The search is guided by an objective function $f : \mathcal{X} \rightarrow \mathbb{R}$, which assigns a scalar value to each solution to measure its quality (or cost). The optimization goal is to find a solution $x^* = \arg \max_{x \in \mathcal{X}} f(x)$ ² [1]. In the context of EAs, this is usually referred to as the “fitness function”, because it measures the viability of an individual solution within its environment.

Generation Mechanism New candidate solutions are created by generation mechanisms, which are stochastic operators designed to navigate the search

²Note that if the objective function represents costs, the function needs to be minimized instead of maximized: $x^* = \arg \min_{x \in \mathcal{X}} f(x)$ or equivalently the function value inverted $x^* = \arg \max_{x \in \mathcal{X}} -f(x)$

space by balancing exploration (diversification) and exploitation (intensification). GAs, for instance, employ two primary mechanisms: crossover (recombination) and mutation. Crossover is a binary operator $(o_1, o_2) = c(p_1, p_2)$ that combines information from two parent solutions, whereas mutation is a unary operator $x' = m(x)$ that introduces small, random changes to a single solution to maintain diversity [50]. This contrasts with simpler mechanisms like the neighborhood perturbation operator in SA [70].

Selection Strategy The selection strategy determines which solutions persist and impact future search directions. As the name suggests, population-based methods like GAs maintain a population at every generation t . Each population is a set of N solutions $P^{(t)} = \{x_1^{(t)}, \dots, x_N^{(t)}\}$. From this set, selection operators choose parents for reproduction, usually based on their fitness. In contrast, single-solution heuristics like SA use an acceptance criterion (e.g., the Metropolis criterion $P(\text{accept}) \propto e^{-\Delta f/T}$ [100]) to decide whether to switch from the current solution to a newly generated one [70].

Initialization and Termination Finally, the entire process is governed by initialization and termination criteria. The algorithm starts with an initial solution or population $P^{(0)}$, often generated uniformly at random from $\mathcal{X}$. The process is stopped based on preset conditions, such as reaching a maximum number of function evaluations, exceeding a computational budget, or detecting convergence of the best-found objective value [54].

2.2. Quantum Computing

QC investigates how information can be encoded in, manipulated and extracted from physical systems that obey the laws of quantum mechanics. Since Feynman’s 1982 proposal to simulate physics with “quantum computers” [44], the field has progressed from theoretical constructs such as the universal quantum Turing machine [35] to early-stage processors with qubits numbers in the order of $\mathcal{O}(10^3)$ [3]. Milestones such as Shor’s polynomial-time factoring algorithm [131] and Grover’s quadratic search speed-up [53] have established hope of asymptotic advantages over classical computation. Today, the current state of NISQ devices is driving research more toward algorithms that tolerate imperfections [118], such as variational algorithms explored in section 2.3.1.

2.2.1. Quantum-Mechanical Principles

Quantum mechanics forms the basis for QC, enabling qubits to behave in ways inherently different from classical bits fixed in states 0 or 1. A single qubit

can exist in a superposition of basis states, while multiple qubits can become entangled, producing correlations that exceed classical logic. Quantum amplitudes can also interfere constructively or destructively based on their complex phases, allowing algorithms to amplify correct outcomes while suppressing incorrect ones [53]. The following presents the quantum-mechanical phenomena that enable quantum computation.

Superposition Superposition is the principle that a quantum system can exist simultaneously in multiple states until it is measured. This means that a single qubit can be in a linear combination of its computational basis states $|0\rangle$ and $|1\rangle$, written as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with complex amplitudes α and β satisfying $|\alpha|^2 + |\beta|^2 = 1$ where $|\alpha|^2$ and $|\beta|^2$ are the probabilities of measuring the qubit in each respective state. The linearity of quantum mechanics allows these combinations to remain in coherent superposition, enabling a register of n qubits to encode amplitudes for all 2^n basis states simultaneously [105]. This exponential state space is the basis of quantum parallelism, where quantum algorithms evaluate functions across many inputs at once within a single step, before interference and measurement select the final outcome.

Entanglement Entanglement is a quantum correlation in which the state of one qubit is inseparably linked to the state of another, regardless of spatial separation. Formally, a multi-qubit system is entangled if its overall state cannot be written as a tensor product of single-qubit states. A canonical example is the Bell state [10]: $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, which shows correlations that cannot be explained by any classical theory with local hidden variables. Entanglement plays a central role in QC: it enables quantum speed-ups [53, 131], robustness in error-correcting codes [45, 51, 133], quantum teleportation [12] and it underlies protocols in secure communication [40]. It is widely regarded as one of the most important resources differentiating quantum from classical information processing.

Interference Interference refers to the way quantum probability amplitudes combine when different computational paths lead to the same outcome. Unlike classical probabilities, which simply add up, quantum amplitudes can add up constructively or destructively depending on their relative phases. This allows quantum algorithms to increase the probability of measuring correct results while suppressing incorrect ones. The use of interference is integral for many quantum speed-ups, including Grover's algorithm [53] for unstructured search and the phase estimation routines underlying Shor's algorithm [131].

Measurement Measurement is the process by which a quantum state is projected onto one of the eigenstates of a specified measurement operator, resulting in a classical outcome. Formally, a measurement is described by a set of




Figure 2.1.: Bloch sphere visualization of a single-qubit state before and after applying a quantum gate.

measurement operators $\{M_m\}$ satisfying $\sum_m M_m^\dagger M_m = \mathbb{I}$, where each M_m corresponds to a possible outcome m . For a single qubit in state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, a projective measurement in the computational basis $\{|0\rangle, |1\rangle\}$ results in 0 with probability $|\alpha|^2$ and 1 with probability $|\beta|^2$. However, measurements can also be performed in other bases, such as the Hadamard basis $\{|+\rangle, |-\rangle\}$ or any arbitrary eigenbasis, depending on the choice of observable. Since measurement outcomes are probabilistic and destroy the quantum state, quantum algorithms must be deliberately designed so that the desired result emerges with high probability when measured in the chosen basis.

2.2.2. Quantum Gates

Quantum gates are the basic operations that transform the state of a quantum system to a subsequent state. Each gate is defined as a unitary matrix acting on one or more qubits, preserving the norm of the state vector. Due to their unitary nature, quantum gates are reversible operations. Single-qubit gates are defined as 2×2 matrices and can be visualized as rotations on the Bloch sphere (see fig. 2.1). Common examples include the Pauli gates X , Y and Z , the Hadamard gate H , which creates superpositions and the phase gates S and T that apply fixed phase shifts. Mathematically, these are given by:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}.$$

Multi-qubit gates act on two or more qubits and can be used to create entanglement and thus can enable quantum computational speedups [110]. The most commonly used entangling gate is the controlled-NOT (CNOT), which flips the state of the target qubit conditioned on the control qubit being in the state $|1\rangle$:

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Other two-qubit gates include controlled-Z [110], iSWAP [123] and the Mølmer-Sørensen XX gate [140], depending on the quantum hardware platform. A gate set is called universal if it can approximate any unitary operation to arbitrary precision. One example of a universal set is the Clifford + T gate set $\{H, S, \text{CNOT}, T\}$, which is sufficient to construct any quantum circuit when combined appropriately [9]. However, current hardware implementations of quantum gates are not ideal; they suffer from various forms of noise such as depolarizing or amplitude-damping channels [110] and coherent over-rotations [56]. These imperfections are quantified using metrics like average gate fidelity [109] and motivate the use of error-correcting codes [45, 133] and fault-tolerant constructions [132] to ensure reliable computation on noisy hardware.

2.2.3. Quantum Circuits

Quantum circuits provide a graphical and mathematical representation of quantum algorithms by arranging quantum gates in a sequence that operates on a register of qubits. Each horizontal wire in the diagram represents the state of a single qubit and gates are applied in order from left to right, as illustrated in fig. 2.2. Formally, a quantum circuit consists of a sequence of layers, where each layer is a tensor product of gates that can be applied in parallel. The overall unitary transformation of the circuit is constructed by composing the unitaries of each layer. An m -gate circuit on n qubits creates a unitary $U = U_m \cdots U_2 U_1 \in SU(2^n)$ followed by a measurement on a chosen subset of qubits. Important structural properties of a circuit include its width (the number of qubits used), depth (the number of sequential gate layers) and total gate count. These parameters influence the algorithm's efficiency and sensitivity to noise, especially in the NISQ era where coherence time is limited.

Quantum circuits must first be transpiled into the specific gate set and topology unique to the device in order to operate on actual hardware. This involves inserting SWAP gates for connectivity constraints [94], decomposing unitaries into native gates [9] and applying optimizations like peephole rewriting or gate fusion [136]. Many circuit identities, such as $HXH = Z$ and Euler decompo-




Figure 2.2.: Example quantum circuit with Hadamard, CNOT and Z gates, followed by a measurement on q_0 .

sitions, help reduce depth and error accumulation. Some devices support mid-circuit measurements [30] and classically controlled operations, which enables error correction, teleportation and adaptive protocols. Logical qubits can be encoded using quantum error correction codes, such as the surface code [45]. This allows for fault-tolerant gates and exponentially suppressed error rates through the use of lattice surgery [61] and transversal gates [38].

2.3. Quantum Artificial Intelligence

QAI investigates how QC can be used to enhance or accelerate AI and, vice versa, how AI methods can support the development and operation of quantum technologies. It represents a rapidly evolving interdisciplinary field that bridges the principles of quantum mechanics and the methodologies of computational intelligence. The idea of combining quantum computation with AI has been explored theoretically for decades, driven by the promise of simulating complex probabilistic systems using quantum mechanics [44, 131]. By harnessing superposition, entanglement and interference, quantum computers realize inherently different computational capabilities than classical machines. This has the potential to enable speedups for tasks that are otherwise infeasible or inefficient. Practical progress, however, remained limited until the emergence of programmable quantum hardware and algorithmic frameworks tailored to NISQ devices [118], which has since sparked renewed interest in quantum-enhanced learning, reasoning and decision-making.

The field of QAI covers a wide range of subareas, most notably QML, which investigates how quantum algorithms, like variational circuits or quantum kernels, can accelerate tasks like classification, clustering and generative modeling. Beyond QML, QAI includes quantum approaches to automated planning [37], natural language processing [155], computer vision [57, 91] and multi-agent systems [39], with applications ranging from logistics optimization and drug discovery to finance and climate modeling [34]. Each of these domains explores if quantum resources can provide computational advantages, improved model expressivity, or better generalization compared to their classical counterparts. In addition to these efforts, AI methods themselves are increasingly used to support the development and operation of quantum technologies. This includes the use of ML for quantum control [111], error mitigation [22], circuit design

[27, 33, 159] and transpilation [117, 147]. Such “AI for quantum” techniques are especially valuable in the current era of NISQ devices, where intelligent adaptation and optimization are essential to deal with hardware limitations.

This thesis covers both directions of QAI: the use of QC to enhance AI and the application of AI techniques to support QC. The latter, commonly referred to as AI for quantum, includes methods for optimizing quantum circuits and the automated design of quantum circuits, and draws on foundations introduced in the preceding sections on classical AI and ML (see section 2.1). The remainder of this chapter, however, focuses on the former direction: quantum algorithms for AI tasks. In particular, it examines variational quantum models, which form the backbone of many near-term QAI approaches. The following sections introduce the core components of these models, including VQCs, input encoding schemes and ansatz design.

2.3.1. Variational Quantum Algorithms

VQAs are a popular class of hybrid quantum-classical algorithms designed to make effective use of current NISQ hardware [23]. They operate by parameterizing a quantum circuit, referred to as a VQC, whose parameters are optimized by a classical computer to minimize a loss function defined over measurement outcomes. The general architecture of a VQA is illustrated in fig. 2.3, where quantum and classical components interact iteratively. This framework enables the use of quantum circuits as expressive, trainable models while leveraging classical optimization to guide the learning process.

VQAs emerged in the mid-2010s as a response to the limitations of QC hardware, with early examples including the Variational Quantum Eigensolver [114] and the Quantum Approximate Optimization Algorithm (QAOA) [42]. These methods demonstrate how useful information can be extracted from shallow quantum circuits by iteratively tuning gate parameters to minimize a loss function³. Since then, VQAs have been widely adopted across QML, quantum chemistry and optimization, becoming an essential algorithm design in the NISQ era.

A VQA consists of three main components: a VQC $U(\theta)$, a classical loss function $\mathcal{L}$ and an optimization algorithm. The circuit prepares a quantum state of n qubits, $|\psi(x, \theta)\rangle = U(x, \theta)|0\rangle^{\otimes n}$, where x denotes input data and θ the trainable circuit parameters. Measurements on this state produce classical outcomes from which the value of $\mathcal{L}$ is estimated⁴. The parameters θ are then updated using a classical optimizer, such as SGD, to minimize the loss. Gradients can be computed using the parameter-shift rule [124], which

³The loss function is in some cases also referred to as a cost function or objective function (see section 2.1.3).

⁴The estimation is commonly obtained through the expectation value of a chosen observable, or through a task-specific loss function, such as classification error.




Figure 2.3.: General structure of a VQA. Blue boxes represent quantum components, while green boxes indicate classical components. The parameters of the circuit are denoted by θ , while θ' refers to the updated parameters produced by the classical optimizer.

enables analytic differentiation of quantum gates by evaluating the circuit at shifted parameter values. This feedback loop continues until convergence or a predefined stopping criterion is met.

The effectiveness of a VQA critically depends on the design of its variational circuit, the encoding of classical data into quantum states and the choice of ansatz. These aspects influence the expressivity of the model, the landscape of the loss function and the feasibility of optimization. The following sections therefore explore each of these components in detail, beginning with the structure and properties of VQCs.

2.3.2. Variational Quantum Circuits

VQCs, also called Parametrized Quantum Circuits (PQCs), are the core component of VQAs. They consist of parametrized quantum gates arranged in a fixed architecture, whose parameters are optimized to minimize a loss function. A VQC can be understood as a quantum analog of a NN. It defines a function f_θ that maps inputs to quantum states, with trainable parameters θ controlling the unitary evolution. The expressive power and trainability of a VQC are determined by its structure, the choice of input encoding and the form of the variational ansatz.

A typical VQC consists of multiple layers of quantum gates applied to an initial state—usually $|0\rangle^{\otimes n}$. These layers include data encoding gates, parametrized single-qubit rotations and entangling operations. The circuit outputs a quantum state from which classical information is extracted via measurement. Due to the unitarity of quantum gates (section 2.2.2), quantum transformations are norm-preserving and maintain the total probability amplitude. Learning is achieved by adjusting gate parameters to exploit constructive and destructive interference in the Hilbert space [23], guided by classical optimization




Figure 2.4.: General structure of a VQC with n qubits and l variational layers. The input data $x = \{x_1, x_2, \dots, x_n\}$ is first embedded into the quantum circuit using an embedding strategy (e.g., Angle Embedding) represented by the blue block. The data is then processed by l variational layers (depicted in green), each consisting of parametrized single-qubit rotations and entangling gates (layer indices are omitted for clarity). Finally, the red block denotes the measurement stage, where the circuit is repeatedly measured to extract classical information.

routines (as discussed in section 2.3.1). The following sections describe two important design choices in constructing VQCs: how classical input data is encoded into quantum states and how the variational ansatz is structured.

Input Embedding The input embedding, also referred to as data encoding or feature mapping, specifies how classical data is translated into quantum states. Since quantum algorithms operate only on quantum states, the first step is to encode an input vector $x \in \mathbb{R}^d$ into a quantum state $|\phi(x)\rangle$. Several encoding strategies exist, each with trade-offs in expressivity, implementation cost and circuit depth. The most common methods include basis encoding, amplitude encoding [105] and angle encoding (depicted in fig. 2.4). Basis encoding maps bits of classical data directly onto computational basis states, but requires many qubits for high-dimensional data. Amplitude encoding represents classical data using the amplitudes of a quantum state, allowing 2^n real values to be stored in the amplitudes of an n -qubit system, thereby offering exponential compression of its representation. However, this advantage depends on the ability to efficiently prepare such quantum states [130], which is usually not trivial in practice. In contrast to amplitude encoding, angle encoding uses data values as rotation angles in parametrized gates (e.g., $R_y(x_i)$) and is widely used in VQCs due to its simplicity and compatibility with hardware-native gate sets. In practice, input encoding is typically integrated into the first layer of the circuit and may be repeated across layers (called data-reuploading [113]) to increase expressivity.

Ansatz The ansatz or variational layer (depicted green in fig. 2.4) defines the structure of the parametrized part of a VQC, i.e., the template of quantum gates with trainable parameters. A well-chosen ansatz is important: it must be expressive enough to approximate the target function or quantum state, but structured enough to allow efficient optimization and implementation. The ansatz determines the hypothesis class that the variational algorithm can explore. Ansätze are often composed of alternating layers of parametrized single-qubit rotations and entangling gates, such as controlled-NOT or Ising-type couplings. Two common design approaches are hardware-efficient ansätze [66], which are tailored to the connectivity and gate set of a specific quantum device and problem-inspired ansätze [42], which embed domain-specific structure into the circuit, e.g., from chemistry or combinatorial optimization. The depth and entanglement structure of the ansatz influence both the expressivity and the optimization landscape. Deep or highly entangled circuits can suffer from barren plateaus, where gradients vanish exponentially with the number of qubits [99]. Its design, possibly incorporating symmetries or adaptive schemes, is therefore essential for practical performance.

3. Training Optimization Techniques

Perfection is achieved, not when there is nothing more to add, but when there is nothing left to take away.

Antoine de Saint-Exupéry, 1939 [122]

The development of QAI models in the NISQ era is inherently limited by the capabilities of available quantum hardware [118]. These devices suffer from noise¹ [71, 118] and a lack of fault-tolerant error correction [148], which collectively hinder scalability. As introduced in sections 2.3.1 and 2.3.2, VQAs offer a practical framework to mitigate these issues by using relatively noise-resilient VQCs within a hybrid quantum-classical optimization loop. Although these architectures present a promising direction, the QAI software ecosystem remains less mature than that of classical AI. In practice, current approaches in the NISQ era still face significant challenges that reduce their applicability. One of the biggest obstacles includes the limited number of qubits available on present-day devices, which severely restricts the complexity of the models. In addition, the training of VQAs is often hindered by barren plateaus, i.e., areas with vanishing gradients that make optimization inefficient or even infeasible [99]. These issues are amplified by instabilities in the hybrid training process, which often leads to poor convergence and unreliable performance across problem instances.

To address these limitations, we introduce several contributions to training optimization techniques in this chapter. First, we develop an approach for weight re-mapping [75, 76], in which the quantum parameters are restricted to their natural periodic range. This approach aligns the optimization with the physical properties of quantum circuits. Second, we propose exponential learning rate decay [87] to improve stability and robustness in variational QRL. This is complemented by a systematic assessment of common heuristics, including data re-uploading [113], input and output scaling [137] and alternative initial-

¹Quantum noise refers to uncontrolled variations in a quantum system that corrupt the stored or processed information. It comes from interactions with the environment (decoherence) and imperfect gate operations, and is a primary obstacle to building reliable quantum computers [118].




Figure 3.1.: The weight re-mapping technique integrates a non-linear function between the classical optimizer and the VQC. This function constrains the optimizer’s unbounded outputs to a physically meaningful angular interval, such as $[-\pi, \pi]$, before they are used to parameterize the quantum gates. Figure adapted from Kölle et al. [75].

ization techniques. Finally, we present a structural pruning method inspired by the Lottery Ticket Hypothesis (LTH). Here, subcircuits are systematically selected to reduce redundancy and improve trainability [78].

Beyond the scope of this chapter, we have made several more training optimization contributions. Gradient-free optimization strategies are investigated as alternatives to gradient-based training, with a particular emphasis on metaheuristic methods (see section 5.2²). To address the constraint of limited qubit counts, we propose a quantum-hybrid architecture in which inputs are compressed prior to quantum processing, extending the applicability of NISQ hardware to more complex tasks such as image and audio classification (see section 6.3³). Finally, we optimize the VQC architecture itself through automated synthesis and circuit optimization techniques, enabling more efficient and scalable model construction in chapter 4.

²Note that we discuss metaheuristic methods in section 5.2 because they are especially important to the training success in QRL.

³Since the emphasis is on the applications and the quantum and classical component contributions, this publication is discussed in section 6.3.

3.1. Adapting Classical Optimization Paradigms

A core problem in training VQCs lies in the mismatch between the optimization tools developed for classical systems and the physical realities of quantum circuits. To address this problem, in Kölle et al. [75, 76] we introduce the method of weight re-mapping, motivated by the inherent inefficiency of applying classical optimizers to quantum parameters. Most classical optimizers operate in an unbounded Euclidean space, whereas the parameters of VQCs are periodic rotational angles. Standard gradient updates that are not aware of this periodicity can cause large, unstable jumps in the parameter space (e.g., from an angle of $\pi - \epsilon$ to $\pi + \epsilon$), which complicate smooth and efficient convergence. The main idea is to enforce this physical constraint directly within the model’s architecture to create a more well-behaved optimization landscape. Our methodology, illustrated in fig. 3.1, involves inserting a non-linear mapping




Figure 3.2.: Training and validation curves of the top 3 approaches for datasets Iris, Abalone, Diabetes and Seeds with Angle Encoding. Figure adapted from Kölle et al. [76].

function between the classical optimizer’s parameter output and the rotational gate’s input. This function constrains the optimizer’s unbounded output to a physically meaningful and non-redundant angular interval, typically $[-\pi, \pi]$, before it is used to parameterize a gate. We systematically evaluate seven different mapping functions (e.g., tanh, arctan and sigmoid) by integrating them into a variational classifier and testing their performance on eight standard classification datasets, using both angle and amplitude embedding.

Our results demonstrate that weight re-mapping consistently accelerates training convergence across all tested datasets (see fig. 3.2). A key contribution is the finding that this simple architectural modification not only speeds up training but, in several cases, also leads to statistically significant improvements in the final test accuracy. This is particularly notable when using angle embedding, where the final accuracy was often higher for models employing a re-mapping function. By creating a more structured optimization landscape, this technique also serves as an effective strategy for mitigating the barren plateau phenomenon [141], a benefit that is explored in section 3.3. Instead of standard QML approaches that treat VQC parameters analogously to classical NN weights, weight re-mapping respects the angle periodicity. While other methods focus on developing entirely new quantum-aware optimizers or complex initialization strategies, we offer a simple, low-overhead solution that is agnostic to the choice of optimizer. It can be seen as an analog to data normalization in classical ML, a simple but highly effective technique for improving model trainability.

3.2. Training Heuristics

Beyond individual parameter adjustments, the overall stability and performance of a VQA depend heavily on the heuristics used [113, 137]. To clarify their effectiveness, in Kölle et al. [87] we conduct a systematic evaluation of established techniques for improving VQC performance, particularly in the sensitive setting of QRL. Within a unified framework using a Quantum Proximal Policy Optimization (QPPO) agent, data re-uploading, learning rates, initialization strategies, input scaling and output scaling are tested in the Frozen Lake and Cart Pole environments (fig. 3.3). This enables a comparative assessment of their individual contributions to overall training dynamics and long-term stability. During this investigation, the strong hyperparameter sensitivity of VQCs in QRL became apparent; agents often require a high learning rate for initial exploration but a much lower one for stable convergence. To address this, we introduce exponential learning rate decay. As shown in fig. 3.4, this schedule provides the optimizer with an appropriately high learning rate at the start of training, which is then systematically reduced according to the decay schedule to allow for stable convergence.




Figure 3.3.: Training architecture of a hybrid quantum-classical actor-critic agent. Evaluated training heuristics are highlighted in green. Figure adapted from Kölle et al. [87].

Our results provide a comprehensive assessment of the evaluated heuristics. The study confirms that data re-uploading (fig. 3.4a) and the newly introduced exponential learning rate decay (fig. 3.4b) enhance hyperparameter stability and overall agent performance. In contrast, input scaling does not improve parameter efficiency, while output scaling proves effective for managing policy greediness, leading to faster and more robust learning. This work’s primary contribution is both an evidence-based comparison of popular training heuristics and the introduction of exponential learning rate decay as a practical and effective strategy for stabilizing VQC training in the notoriously unstable RL context.

3.3. Strategies for Mitigating Barren Plateaus

The barren plateau phenomenon—where the optimization landscape of a VQA becomes exponentially flat and featureless, causing gradients to vanish as system size increases—remains one of the most significant barriers to scaling VQAs. In this section, we present several methods to tackle this problem. Although not demonstrated in this thesis, prior work suggests that the weight re-mapping technique introduced in section 3.1 can also help reduce barren plateaus. In particular, Stenzel et al. [141] provide evidence that constraining the effective parameter space through re-mapping counteracts the exponential flattening of the loss landscape by keeping gradients more pronounced.




Figure 3.4.: Performance of the training heuristics introduced in fig. 3.3 for FrozenLake, along with a detailed comparison of learning-rate schedules in CartPole. Figure adapted from Kölle et al. [87].

On this basis, we investigate two additional strategies for mitigating barren plateaus: gradient-free optimization methods that circumvent the vanishing gradient problem directly⁴ and model simplification techniques that reduce the complexity of the optimization landscape. In particular, in Kölle et al. [84] we explore the integration of various metaheuristic algorithms into QRL frameworks, which constitute a central component of the approach in this thesis and are discussed in more detail in section 5.2. Another effective strategy is to reduce the complexity of the VQC itself. Our work in Kölle et al. [78] is motivated by the idea that over-parameterization is a key contributor to barren plateaus. We apply the LTH from classical ML to the quantum domain, proposing that a large, dense VQC contains a smaller, trainable subcircuit (a “winning ticket”) capable of matching its performance, thereby simplifying the optimization landscape. The weak LTH states that after training and pruning, there exists a sparse subnetwork which, when reset to its original initialization and retrained, can match the performance of the full model, while the strong LTH claims that such subnetworks can be found without any training by learning only a pruning mask. To evaluate this, the weak LTH is tested using a train–prune–retrain methodology: a full, over-parameterized VQC is trained, low-magnitude weights are pruned and the remaining sparse subcircuit is reset to its original initialization and retrained. Both iterative and one-shot pruning strategies are applied to two VQC architectures and a classical NN baseline on variants of the Iris and Wine classification datasets. In contrast,

⁴Note that gradient-free approaches do not eliminate the barren plateau problem, as the cost function remains unchanged. Instead, they avoid vanishing gradients by relying on optimization strategies not prone to this issue, such as population-based algorithms, which can explore a larger search space at the expense of increased computational resources. Nevertheless, barren plateaus can still pose a challenge if the optimization landscape is so large and flat that even population-based methods cannot sufficiently explore it.

the strong LTH is investigated with a custom EA that directly searches for effective pruning masks without any gradient-based training.




Figure 3.5.: Performance comparison of a pruned VQC (winning ticket, orange line) against its unpruned counterpart (blue line) on the Wine dataset. The pruned model, retaining only 32.7% of its weights, significantly outperforms the unpruned model, demonstrating that pruning can be an effective strategy to escape a barren plateau and improve generalization. Figure adapted from Kölle et al. [78].

Table 3.1.: Summary of winning tickets, showing the percentage of remaining weights required for a pruned subcircuit to match or exceed the performance of the unpruned model across different datasets and VQC architectures. BVQC denotes a binary-classifier VQC, MVQC a multi-class-classifier VQC and SNN a simple multi-class NN. Table adapted from Kölle et al. [78].

Dataset	BVQC	MVQC	SNN
Simplified Iris	33.3%	26.1%	20.8%
Iris		26.0%	32.1%
Simplified Wine	51.3%	-	51.4%
Wine		32.7%	21.1%

Our results confirm that the weak LTH holds for VQCs, identifying winning tickets with as few as 26% of the original parameters while matching or exceeding the full model’s performance, as summarized in table 3.1. The most significant contribution is the demonstration that pruning can serve as a powerful optimization technique to escape barren plateaus. As illustrated in fig. 3.5, on the Wine dataset, a pruned VQC (fig. 3.5a) achieved 80% accuracy, drastically outperforming the unpruned model’s 45% accuracy, which

was presumably trapped in a barren plateau. Importantly, we observe this improvement also when compared against classical sparse NNs (fig. 3.5b), suggesting that the benefit is not only due to reduced model complexity. Our approach contrasts with other methods for mitigating barren plateaus that focus on initialization strategies, local cost functions, or alternative optimizers. By formulating the problem as one of structural redundancy, the LTH provides a novel, architecture-focused solution. It connects to the broader field of NN pruning but is one of the first works to systematically apply this hypothesis to VQCs and explicitly demonstrate its utility in overcoming a failed training process.

4. Automated Circuit Design and Synthesis

What I cannot create, I do not understand.

Richard P. Feynman, 1988 [48]

As detailed in section 2.3.1, the VQC is the core of a VQA. Its architecture determines both the circuit’s expressiveness and its trainability, making it decisive for the overall performance. However, designing suitable architectures is often a challenging and labor-intensive task [23]. Traditionally, the field has relied on manual, heuristic methods [13] to construct ansätze. These are typically either hardware-efficient, designed to minimize costly gates and connectivity for feasibility on specific quantum devices [66, 93], or problem-inspired, embedding knowledge of the problem’s underlying physics directly into the circuit [98, 114]. While these approaches provide useful starting points, they are inherently constrained by human intuition and often require substantial domain expertise [52]. This dependence on manual design therefore represents a bottleneck, limiting the scalability and broader applicability of QAI. Automated design offers a way to systematically explore the extensive combinatorial space of possible circuits [158]. This paradigm shift, similar to how DeepMind’s AlphaTensor [43] discovered novel matrix multiplication algorithms, formulates circuit design as a search problem to be solved by intelligent algorithms [158]. By automating this discovery process, non-intuitive yet highly efficient architectures can potentially be uncovered that, in DeepMind’s case, even surpassed their human-designed counterparts. In this chapter, we investigate two powerful approaches for this task: RL, where an agent learns a sequential policy to construct circuits gate by gate, and EAs, which apply principles of natural selection to evolve populations of circuits toward higher performance. Building on the methods for training a given circuit established in chapter 3, this chapter focuses on finding the right circuit to train.

4.1. State Preparation using Reinforcement Learning

A central challenge in many areas of quantum research is the preparation of specific quantum states, where ML has emerged as a promising tool for automation [88, 89, 96]. Recently, Gabor et al. [46] developed an RL approach to generate unitary transformation matrices, but their method’s reliance on a costly QR-decomposition and the use of density matrices limits its efficiency. In Kölle et al. [83], we develop a new RL framework for automated quantum circuit design that improves efficiency by using only unitary transformations and a computationally streamlined state vector representation of the quantum state, at the cost of accuracy. The motivation is to move beyond manual and often ad-hoc circuit design and to establish a systematic, automated method for discovering the sequence of quantum gates required to prepare a specific target quantum state. This is a foundational task in QC [130] and automating it provides a building block for more complex algorithmic design. We develop a comprehensive and open-source RL environment where an agent learns to construct a quantum circuit. The environment is formulated as an MDP: the state space consists of a vector representation of the current quantum state and the target state. The action space is discrete and includes the set of universal gates (the Clifford+T set) and all possible qubit locations to which they can be applied. At each step, the agent selects a gate and a location, applies it to the current state and receives a reward. The reward function is designed to guide the agent toward the target state with high fidelity while penalizing excessive circuit length. A PPO agent is used to learn a policy for this sequential construction task.

Table 4.1.: The 2-qubit benchmark set, showing the probability that the best-performing agents successfully generated the target quantum state. The high success rates demonstrate the agent’s ability to find effective circuit designs for a variety of states. Table adapted from Kölle et al. [83].

State	Gen. probability [%]
$\frac{1}{2}(00\rangle + 01\rangle + 10\rangle + 11\rangle)$	100
$\frac{1}{\sqrt{2}}(00\rangle + 11\rangle)$	100
$ 01\rangle$	96
$ 10\rangle$	98
$\frac{1}{\sqrt{2}}(00\rangle - 11\rangle)$	76
$ 11\rangle$	99
$\frac{1}{\sqrt{2}}(01\rangle + 10\rangle)$	96

The primary contribution of this work is the development of an integrated and scalable RL environment for quantum circuit synthesis, as shown in fig. 4.1.




Figure 4.1.: Performance of RL agents on state preparation tasks for varying qubit numbers (n) and circuit depths (λ). Each data point shows the average performance of three agents, demonstrating the framework’s modularity to handle different system sizes. Figure adapted from Kölle et al. [83].

Our results demonstrate that a trained PPO agent can consistently find circuits with minimum circuit depth that correspond to creating common 2-qubit states, such as the Bell states. The discovered circuits match the efficiency of known, manually designed optimal solutions, as detailed in table 4.1. This validates the RL approach as an effective method for automating small circuit synthesis tasks. However, applying the method to circuits with more qubits or greater depths remains difficult and will require further improvements. Our approach is fully automated, data-driven, and differs from other ML approaches that generate potentially non-unitary matrices and require costly post-processing (e.g., QR-decomposition [46]) to ensure physical validity. By operating directly with a discrete, universal set of unitary quantum gates, this method guarantees by construction that the synthesized circuits are valid and can be implemented on quantum hardware. It serves as a foundational step from which more complex architecture search methods can be built.

4.2. Ansatz Design using Reinforcement Learning

In Kölle et al. [81], we extend the scope from simple state preparation to the more complex task of designing VQC ansätze for specific QML problems. Since the performance of a VQC is highly dependent on its architecture, or ansatz, manually creating an effective design that is both expressive and trainable remains a key challenge in QML. This work addresses the need to automate this design process, allowing an agent to discover architectures tailored to a specific




Figure 4.2.: The overall architecture of the RL-QAS framework. The outer loop, managed by an RL agent, is responsible for constructing candidate VQC architectures. The inner loop evaluates the performance of these architectures by training them as a VQC on a given task. Figure adapted from Kölle et al. [81].

ML task, improving performance and reducing the need for manual heuristics by creating a Quantum Architecture Search (QAS) framework driven by RL. We introduce an RL-QAS framework, depicted in fig. 4.2, which employs a two-loop structure. In the outer loop, an RL agent constructs a candidate VQC architecture by sequentially placing gates. The state is represented by a 3D binary tensor that encodes the structure of the circuit similar to Patel et al. [112]. The action space is multidiscrete, where gate types and qubit indices can be selected. In the inner loop, the constructed circuit is used as the ansatz within a VQC, which is then trained and evaluated on a specific QML classification task (e.g., Iris or binary Modified National Institute of Standards and Technology (MNIST)). The performance on the task (i.e., test accuracy) is used to calculate a reward for the RL agent, which is also penalized for circuit complexity (gate count and depth). This reward guides the agent in the outer loop to discover better architectures over time.

The main contribution is a complete, end-to-end RL framework for QAS in a QML context. Our results, summarized in table 4.2, show that the agent autonomously discovers low-complexity circuit designs that achieve high test accuracy on the Iris and binary MNIST datasets. The discovered architectures consistently achieve comparable or superior performance to hardware-efficient, manually designed ansätze, such as strongly entangling layers [125], while using significantly fewer gates. This demonstrates that RL is a viable and power-

Table 4.2.: Comparison of VQCs designed by RL-QAS versus manually designed Strongly Entangling Layer (SEL) [125] ansätze across five datasets. The metrics include gates, parameters, CNOTs, circuit depth, training accuracy (TrA) and test accuracy (TeA). The RL-QAS circuits consistently achieve comparable or superior performance with lower complexity. Table adapted from Kölle et al. [81].


Dataset	Approach	Gates	Param.	CNOT	Depth	TrA	TeA
Iris 2 (0,1)	RL-QAS VQC	1	1	0	1	1.0	1.0
	SEL VQC (1)	8	6	2	4	0.98	1.0
Iris 2 (0,2)	RL-QAS VQC	1	1	0	1	1.0	1.0
	SEL VQC (1)	8	6	2	4	0.98	1.0
Iris 2 (1,2)	RL-QAS VQC	3	2	1	2	1.0	1.0
	SEL VQC (1)	8	6	2	4	0.98	1.0
Iris	RL-QAS VQC	4	3	1	3	0.95	1.0
	SEL VQC (1)	8	6	2	4	0.66	0.66
	SEL VQC (2)	16	12	4	8	0.84	0.82
MNIST2	RL-QAS VQC	14	6	8	7	0.84	0.91
	SEL VQC (1)	20	15	5	8	0.77	0.77
	SEL VQC (2)	40	30	10	16	0.86	0.93

ful approach for automated architecture search in QML, but again reveals challenges when extending to more complex and deeper circuits. This work represents a significant step up in complexity from the discrete state preparation task in the previous section, moving into the parameterized domain of VQC for ML. Different from other QAS methods like EAs [47] or differentiable search [26, 156], our work uses an RL agent, which can learn a more structured construction policy. The two-loop approach is a well-established pattern in classical NN Architecture Search and this work successfully ports and validates this approach in the quantum domain.

4.3. Ansatz Design using Evolutionary Strategies

In Kölle et al. [82], we explore EAs as an alternative paradigm for automated circuit design, focusing on the co-evolution of both circuit architecture and parameters in a complex Quantum Multi-Agent Reinforcement Learning (QMARL) setting. The primary motivation is to overcome two major hurdles in QRL: the barren plateau problem, which plagues gradient-based training and the reliance on fixed, manually designed circuit architectures. The core idea is to use a gradient-free, population-based EA that can simultaneously optimize a VQC’s parameters and evolve its underlying structure, allowing for a more flexible and robust search for high-performing agents. We propose and compare three distinct evolutionary strategies for representing and evolving

circuit architectures: a ‘Layer-Based’ approach, where entire layers of gates are treated as genetic blocks for mutation and recombination; a more granular ‘Gate-Based’ approach, where individual gates are manipulated; and a hybrid ‘Prototype-Based’ approach, where a prototype (template) layer is built out of individual gates. A population of VQC agents, each with a potentially different architecture, is then evaluated in the cooperative Coin Game environment. The fitness of each agent is determined by its performance (score) in the game. The GA then uses selection, crossover and mutation operators tailored to these representations to produce the next generation of circuits.



(a) Best and average score for the different architectural approaches. (b) Best and average score of Gate-Based approach and baseline approach.

Figure 4.3.: Comparison of the performance of different evolutionary strategies for architectural design in (a). Comparison showing that the flexible, gate-based evolutionary approach outperforms a baseline, non-evolving architecture [125] in (b). Figure adapted from Kölle et al. [82].

The main contribution is the introduction and systematic comparison of these different architectural evolution strategies within an EA framework. Our results, shown in fig. 4.3, reveal several insights: Strategies based solely on mutation are more effective than combining mutation with recombination for this task. The flexible ‘Gate-Based’ approach (fig. 4.3a) significantly outperforms the more rigid ‘Layer-Based’ and ‘Prototype-Based’ methods. Agents are discovered that achieve higher scores with more compact circuits and even outperform a static (non-evolving) architecture (fig. 4.3b). This clearly demonstrates the benefit of allowing the architecture itself to evolve. Unlike many applications of GA in QC that focus solely on parameter optimization, our method allows the circuit structure to be part of the evolutionary search. The finding that flexible, gate-level manipulation is superior to layer-level manipulation provides important guidance for designing future automated design frameworks, suggesting that fine-grained control over the architecture is key to discovering novel and efficient solutions.

4.4. Mutation Strategies for Circuit Design

In Kölle et al. [73], we provide a deeper investigation into an important component of the GA approach: the mutation operator. The performance of a GA is highly dependent on how it explores the solution space and mutation is the main mechanism for adding new variation. This work addresses the fact that while GAs have been proposed for circuit design, there has been little systematic analysis of which mutation strategies are most effective for manipulating quantum circuits. Understanding this is crucial for designing efficient and powerful EAs for this domain. We create a comprehensive GA framework to systematically test different mutation strategies and their combinations. The operators include: ‘change’ (replace a gate with another), ‘delete’ (remove a gate), ‘add’ (insert a new gate) and ‘swap’ (reorder two gates). The fitness function is designed to balance three objectives: maximizing the fidelity to a target quantum state while minimizing both the circuit depth and the count of resource-intensive T-gates. The experiments are conducted on the task of synthesizing circuits for randomly generated 4- and 6-qubit target states.




Figure 4.4.: A histogram of the average final fitness scores for various mutation strategy combinations on 6-qubit state preparation tasks. The results highlight which combinations effectively balance fidelity and resource efficiency. Figure adapted from Kölle et al. [73].

The key contribution is a detailed, empirical analysis and comparison of mutation operators for quantum circuit synthesis. Our results, illustrated in fig. 4.4, reveal that a combined strategy of ‘swap, addition, deletion’ offers the best balance between performance and resource efficiency, delivering near-optimal fidelity while also having the potential to reduce circuit depth. In contrast,

the ‘swap, addition’ strategy achieves the highest fidelity scores but consistently does so at the cost of increasing circuit depth. This work provides concrete, actionable guidance for designing more effective GAs for quantum circuit optimization. This work provides a deep dive into the mechanics of the EA approach presented in the previous section. While many papers on evolutionary quantum circuit design exist, they often use a fixed set of operators without a comparative study. This paper provides that missing analysis and contributes knowledge to the field. It contrasts with methods that focus on the high-level structure of the algorithm (like island models) by instead focusing on the low-level operators that drive the evolutionary search, showing that the choice of these operators can have a significant impact on the outcome.

5. Quantum Reinforcement Learning

[...] in theory there is no difference between theory and practice, while in practice there is?

Benjamin Brewster, 1882 [18]

The integration of QC into AI applications has the potential to address fundamental challenges in computational complexity. By utilizing principles like superposition and entanglement, quantum algorithms may offer significant advantages in time and space efficiency for certain problems [15]. A particularly promising domain for quantum enhancement is RL, a paradigm where an agent learns to make optimal decisions through trial-and-error by interacting with an environment to maximize a cumulative reward signal, as detailed in section 2.1.2. RL is often characterized by computationally intensive tasks, such as navigating vast state-action spaces [143] and learning from high-dimensional inputs [103]. QC has the potential to improve RL by providing more powerful and parameter-efficient function approximators. The current state of QRL is an active and rapidly developing area of research [25], where early results suggest that replacing classical NNs with more compact VQCs enables agents to learn complex policies with significantly fewer parameters [29]. In this chapter, we explore the practical application and further development of QRL. Building on the foundational training heuristics evaluated in chapter 3 (e.g., in Kölle et al. [87]), which provide the necessary tools for stable optimization, three core contributions are explored in the following sections. First, we examine the development of hybrid quantum-classical actor-critic agents, which represents an important step toward improving performance and stability in continuous control tasks. Second, we investigate the use of gradient-free metaheuristic strategies as a direct and practical solution to overcome training obstacles such as barren plateaus. Finally, we extend QRL to the multi-agent domain, allowing quantum models to be tested with regard to scalable, parameter-efficient learning in complex, distributed systems.

5.1. Hybrid Quantum-Classical Actor-Critic Agents

In Kölle et al. [77], we investigate the integration of VQCs into the A2C algorithm to address the scalability concerns of classical RL, where the size of NNs can become a bottleneck. This work tests whether the parameter efficiency of VQCs can be leveraged to create more compact yet high-performing agents for continuous control tasks by replacing classical actor and critic networks with quantum counterparts. The study systematically replaces the actor network, the critic network, or both within a classical A2C framework with VQCs. To address the limitations of using a pure quantum output, a hybrid architecture is also proposed and tested. In this hybrid model, as illustrated in fig. 5.1, the measurement outcomes from all qubits of the VQC are fed into a classical linear-layer for post-processing¹. This architecture can be configured as a hybrid actor (fig. 5.1a) or hybrid critic (fig. 5.1b), allowing the model to scale the quantum output to the desired size (e.g., two actions for the actor, one value for the critic). All configurations are empirically evaluated in the continuous state space of the Cart Pole environment.




Figure 5.1.: The hybrid VQC architecture used for the actor (a) and critic (b). The output of the variational circuit is passed to a classical post-processing NN. Figure adapted from [77].

¹This post-processing layer is not to be confused with a Dressed Quantum Circuit (DQC), where a quantum circuit is sandwiched between two classical NNs, or a transfer learning approach discussed in section 6.3. This methodology is more similar to output scaling introduced by Skolik et al. [137] or Jerbi et al. [64].

The results reveal that the pure quantum A2C configurations fail to learn effectively, due to barren plateaus [99]. The central contribution of this work is the finding that the hybrid quantum-classical strategy results in a substantial performance increase, as shown in fig. 5.2. Both the hybrid actor and the hybrid critic configurations significantly outperform not only the pure quantum variants but also the purely classical A2C algorithm that has a comparable number of trainable parameters. This work contrasts with purely theoretical QRL proposals by emphasizing the practical necessity of classical post-processing in the NISQ era to compensate for the limited output space of quantum circuits. The success of the hybrid model ties in with the overarching theme of hybrid quantum-classical computing and demonstrates a practical way to use the expressive power of quantum circuits while relying on classical layers to process their outputs into a form that is stable and useful for the RL algorithm.




Figure 5.2.: Performance comparison of the hybrid A2C (HA2C) configurations against the classical A2C baseline in the Cart Pole environment. The hybrid approaches demonstrate superior learning. Figure adapted from [77].

5.2. Metaheuristic Optimization for QRL Agents

In Kölle et al. [84], we address the challenge of training VQC-based agents in the presence of flat and deceptive optimization landscapes. This work seeks to find a robust alternative to gradient-based optimization methods, as they are often inefficient for VQCs due to the barren plateau phenomenon, by systematically investigating and comparing a variety of gradient-free, metaheuristic algorithms as optimizers for QRL agents. The study evaluates six different metaheuristic algorithms in particular: Particle Swarm Optimization (PSO), SA, Ant Colony Optimization (ACO), TS, GA, and Harmony Search (HS). Each algorithm is used to optimize the parameters of a VQC agent. Their

performance is evaluated on two different RL environments: a discrete 5×5 MiniGrid navigation task and the continuous Cart Pole control task. The algorithms are compared based on performance, convergence speed and stability.

All tested algorithms in the MiniGrid environment successfully guide the agent to near-optimal results, with SA and PSO performing best (fig. 5.3a). In the more complex Cart Pole environment (fig. 5.3b), SA, GA, and PSO again achieve optimal results, while the other methods perform only slightly better than random, as shown in fig. 5.3. The key contribution of this work is a comparative study showing that metaheuristic optimization is a highly effective strategy for training QRL agents. It provides clear evidence for the potential of PSO and SA as efficient optimizers in QRL and underscores the importance of careful algorithm selection. This approach offers an alternative to the majority of the QRL literature that relies on gradient-based optimizers like Adam. It builds upon the initial promise of evolutionary methods shown in works like [28], but significantly expands the scope by evaluating a much wider range of metaheuristics. It provides a practical, gradient-free solution to the barren plateau problem discussed in chapter 3, establishing a validated set of alternative optimizers for QRL.




Figure 5.3.: Performance of the six metaheuristic algorithms in the 5×5 MiniGrid (a) and Cart Pole (b) environments. PSO and SA consistently emerge as top performers. Figure adapted from [84].

5.3. Evolutionary Multi-Agent Quantum Reinforcement Learning

In Kölle et al. [86], we extend the application of QRL from single-agent settings to the more complex and scalable domain of Multi-Agent Reinforcement Learning (MARL). Multi-agent systems often suffer from an exponential growth in the joint state-action space [21], demanding large and complex models for each

agent. This work leverages the potential of VQCs to mitigate this scalability challenge through their inherent parameter efficiency by demonstrating that a team of agents represented by compact VQCs can learn to solve a cooperative task as effectively as classical agents that require vastly more parameters. Our study uses a gradient-free EA to train a population of VQC-based agents in a multi-agent setting. Each agent’s policy is represented by a VQC. Instead of using gradients, the agents’ collective performance (the total score in the cooperative Coin Game environment) serves as the fitness function. The EA uses selection and mutation to iteratively improve the population of agents over generations. The performance of the VQC agents is compared against two classical baselines: a NN with a similar number of parameters and a much larger NN.




Figure 5.4.: Average score of the VQC agent compared to a small NN with a similar parameter count and a much larger NN. The VQC performs comparably to the large NN with significantly fewer parameters. Figure adapted from [86].

Our results, shown in fig. 5.4, demonstrate that the VQC agents significantly outperform the classical NN of a similar size. The VQC agents achieve a performance comparable to that of the large NNs while using 97.88% fewer parameters. The main contribution is the clear demonstration of a massive parameter-efficiency advantage for VQCs in a multi-agent context. This shows that the benefits of quantum models are not confined to single-agent tasks but can be a powerful tool for creating more scalable MARL systems. This work extends the single-agent, gradient-free QRL paradigm to the multi-agent domain. The use of an EA connects it to the metaheuristic approaches discussed in the previous section, confirming the power of gradient-free optimization in a more complex setting.

6. Advanced Models and Applications

We are not students of some subject matter, but students of problems. And problems may cut right across the borders of any subject matter or discipline.

Karl Popper, 1962 [115]

While the previous chapter focused on the specific domain of QRL, this chapter broadens the scope to examine practical applications of VQAs across diverse ML paradigms. The central motivation is to explore the potential of quantum-enhanced models in modern AI applications, investigating whether their unique properties, as discussed in section 2.2, can be effectively leveraged in more advanced scenarios. In this chapter, we present four contributions across different QML applications. First, we explore quantum-enhanced generative modeling through diffusion models. Although these models represent the state of the art in generative AI for image synthesis [8, 157], their performance comes at a high computational cost, characterized by large NNs and a slow sampling process requiring numerous iterative steps [59, 108]. In the first section, we propose quantum-native diffusion architectures that leverage the unitarity of quantum circuits for more efficient generation. Second, we investigate quantum kernel methods for anomaly detection, where such methods show promise by capturing subtle correlations in high-dimensional spaces [55]. However, prior work like Kyriienko and Magnusson [90] faced quadratic scaling limitations, which we address using efficiency techniques that reduce complexity from quadratic to approximately linear scaling. Third, we present a methodological contribution by addressing a central issue in hybrid quantum–classical models: the frequent use of classical pre- and post-processing—known as DQCs [14, 97]—can obscure the actual contribution of the quantum components to performance. To address this, we develop a transparent architecture in which classical compression techniques are applied for dimensionality reduction. This architecture enables quantum methods to be used with more complex datasets, such as audio and images, while ensuring that the effects of the quantum model are clearly distinguished from those of classical processing. Finally, we conclude the chapter by addressing a foundational claim in QML: that VQCs can out-




Figure 6.1.: A visual comparison of the iterative denoising process for the proposed Q-Dense model against classical and other quantum counterparts on MNIST digits, illustrating the model’s ability to progressively generate coherent images from noise. Figure adapted from Kölle et al. [85].

perform classical NNs with significantly fewer parameters. This promise is rooted in the high theoretical expressivity of quantum models [135] and has been supported by targeted studies [2] and experiments in this thesis. However, the practical trade-offs of this parameter efficiency, particularly when weighed against training times on real quantum hardware, remain largely unquantified in systematic, head-to-head comparisons. The final section of this chapter addresses this gap, providing data-driven analysis of the real-world advantages and current limitations of VQCs.

6.1. Quantum Denoising Diffusion Models

Venturing into the domain of generative AI, in Kölle et al. [85] we explore the integration of VQCs into the architecture of denoising diffusion models, a state-of-the-art class of generative models [8, 157]. Classical diffusion models, while powerful, face two significant challenges: they are often computationally intensive, requiring many sampling steps to generate an image [59, 138] and their underlying NNs can be extremely large and parameter-heavy [119]. This work leverages the potential of QC to address both issues by replacing the classical U-Net¹, the standard workhorse of diffusion models, with more compact and potentially more powerful VQC-based architectures. We introduce two novel quantum architectures: a Q-Dense model, which uses a strongly en-

¹A U-Net is a convolutional neural network architecture originally developed for biomedical image segmentation. Its characteristic U-shaped design, consisting of a contracting path to capture context and a symmetric expanding path for precise localization, has made it a standard for many image-to-image tasks.

tangling VQC [125] as a direct replacement for a dense classical network and a more sophisticated Quantum U-Net (QU-Net), which incorporates a novel quantum convolution layer. Furthermore, a “unitary single-sampling” architecture is proposed. This method uses the inherent unitarity of quantum circuits to concatenate the entire iterative denoising process into a single unitary matrix, allowing for one-shot image generation analogous to classical consistency models [139]. These models are benchmarked against classical counterparts on the MNIST, Fashion MNIST and Canadian Institute for Advanced Research (CIFAR)-10 datasets.




Figure 6.2.: (a) An illustration of the QU-Net architecture, which incorporates novel quantum convolutional layers. (b) A comparison of the FID score for the Q-Dense model and its classical counterparts, demonstrating competitive performance. Figure adapted from Kölle et al. [85].

The primary contribution is the successful design and implementation of the first quantum-native U-Net, as depicted in fig. 6.2a, and a unitary consistency model for diffusion-based image generation. Our results, visualized in fig. 6.1, show that both the Q-Dense and QU-Net models surpass the performance of classical models with similar parameter counts across all metrics (FID, Structural Similarity Index Measure (SSIM), and Peak Signal-to-Noise Ratio (PSNR)), as demonstrated in fig. 6.2b. The unitary single-sampling model is also successfully demonstrated on both simulators and real IBMQ hardware, proving its viability. This work marks a notable advance in the application of QAI to state-of-the-art generative modeling. Unlike much of the existing QML literature, which has primarily addressed discriminative tasks, this contribution extends quantum approaches to the challenging domain of image synthesis. While earlier efforts in quantum diffusion models exist [68], the architectures introduced here are more sophisticated, featuring quantum-




Figure 6.3.: Training and testing time in seconds as a function of training data size, showing the computational efficiency gains of the proposed methods (randomized measurements and variable subsampling) over the standard inversion test (IT). Figure adapted from Kölle et al. [72].

native U-Net designs and unitary consistency models. These findings indicate that the unique properties of VQCs can be effectively harnessed for complex generative tasks, opening new avenues for quantum-enhanced generative AI.

6.2. Efficient Quantum Anomaly Detection

In Kölle et al. [72], we address an important scalability bottleneck in the application of quantum kernel methods to anomaly detection. Quantum kernels, particularly those based on feature maps that are hard to simulate classically [55, 95, 124], have shown promise for enhancing kernel-based algorithms like Support Vector Machines (SVMs). However, their practical application is severely limited by the computational cost of constructing the kernel matrix, which scales quadratically with the number of data points [16, 124, 126]. This work addresses the need to make quantum kernel methods for anomaly detection computationally feasible for real-world datasets. The study focuses on an One-Class Support Vector Machine (OC-SVM) using a quantum kernel. To mitigate the quadratic time complexity, we investigate two distinct efficiency-enhancing techniques. The first is the use of randomized measure-




Figure 6.4.: Performance metrics for the different anomaly detection methods on the Credit Card Fraud dataset, evaluated by the number of features (qubits). The results show that the efficiency-enhancing techniques maintain competitive performance. Figure adapted from Kölle et al. [72].




Figure 6.5.: The proposed hybrid architecture, which consists of a classical Auto-Encoder (AE) for unsupervised dimensionality reduction and a VQC for classification. This design allows for a clear separation of tasks between the classical and quantum components. Figure adapted from Kölle et al. [80].

ments to estimate kernel entries [63], a quantum-native approach that reduces the number of required circuit executions. The second is an adaptation of a classical ensemble method, variable subsampling [36], where multiple OC-SVMs are trained on smaller, random subsets of the data and their results are aggregated. Both methods are evaluated on synthetic data and a real-world credit card fraud detection dataset.

The key contribution is the successful application and validation of two methods that reduce the time complexity of quantum OC-SVMs from quadratic to approximately linear with respect to data size. As shown in fig. 6.3, the experimental results demonstrate a substantial reduction in training and inference times—up to 95% and 25%, respectively. While the randomized measurements approach proves somewhat unstable, the variable subsampling method is able to maintain the performance of the full quantum kernel, as seen in fig. 6.4, while providing significant speedups. This work directly tackles a well-known limitation of quantum kernel methods. By adapting both a quantum-native (randomized measurements) and a classical (variable subsampling) technique, the paper provides a practical comparison of different paths toward making these powerful models more efficient and applicable to real-world anomaly detection problems.

6.3. Disentangling Hybrid Quantum-Classical Models

In Kölle et al. [80] and its extended version Kölle et al. [79], we provide a systematic analysis of hybrid quantum-classical models, aiming to clarify the

true source of their performance. Hybrid models that use a classical NN to compress high-dimensional data before feeding it to a VQC have become a popular strategy for applying quantum circuits to large-scale problems. However, it is often unclear whether the model’s performance comes from the quantum component or if the classical pre-processing network is doing most of the classification work. This work addresses the need for a more transparent methodology to disentangle the respective contributions of the classical and quantum parts. We propose a hybrid architecture, which is illustrated in fig. 6.5. Instead of using a pre-trained classification network for feature extraction (as is common in transfer learning approaches), this model uses a classical AE for dimensionality reduction. The AE is trained only on a reconstruction loss, meaning it is not explicitly trained to classify. The compressed latent space representation from the AE’s encoder is then passed to a VQC for the actual classification task. This design makes it possible to apply quantum classifiers to complex datasets while isolating their performance from the classical pre-processing. The architecture is benchmarked against VQCs in combination with Principal Component Analysis (PCA) and Gaussian Mixture Models (GMMs), as well as state-of-the-art hybrid transfer learning models (DQC [14, 97], SEQUENT [4]), purely classical models and a purely quantum model using amplitude embedding.

Table 6.1.: Comparative test accuracies for all models across four datasets, highlighting the performance of the proposed AE+VQC model against various hybrid and classical baselines. Ang. denotes Angle Encoding, Amp. Amplitude Encoding and comb. indicates that the models were trained combined instead of separate. Table adapted from Kölle et al. [79].

Model	Banknote	Breast Cancer	MNIST	AudioMNIST	Average
AE+VQC (Ang.)	0.79 ± 0.02	0.82 ± 0.11	0.51 ± 0.04	0.24 ± 0.03	0.59 ± 0.43
VQC (Amp.)	0.85 ± 0.04	0.85 ± 0.02	0.44 ± 0.03	0.26 ± 0.03	0.60 ± 0.47
DQC	0.99 ± 0.00	0.97 ± 0.02	0.90 ± 0.01	0.36 ± 0.19	0.81 ± 0.48
SEQUENT [4]	0.98 ± 0.01	0.97 ± 0.02	0.51 ± 0.05	0.39 ± 0.04	0.71 ± 0.49
GMM+VQC	0.58 ± 0.05	0.93 ± 0.02	0.40 ± 0.03	0.12 ± 0.02	0.51 ± 0.54
PCA+VQC	0.79 ± 0.02	0.90 ± 0.03	0.12 ± 0.00	0.23 ± 0.04	0.51 ± 0.62
AE+VQC (comb.)	0.79 ± 0.03	0.79 ± 0.12	0.48 ± 0.03	0.21 ± 0.04	0.57 ± 0.44
AE+NN	0.70 ± 0.08	0.83 ± 0.12	0.83 ± 0.04	0.30 ± 0.03	0.67 ± 0.40
NN	0.99 ± 0.01	0.97 ± 0.02	0.99 ± 0.00	0.88 ± 0.04	0.96 ± 0.08
AE	0.17 ± 0.03	0.02 ± 0.02	$< 0.01 \pm < 0.01$	$< 0.01 \pm < 0.01$	0.05 ± 0.13

Our main contribution is a clear demonstration that classical pre-processing components in popular hybrid models play a significant and often dominant role in classification performance. The analysis, summarized in table 6.1, reveals that the high performance of models like DQC and SEQUENT is largely attributable to their classical layers. The proposed AE-VQC model, where the classification task is explicitly isolated to the quantum component, performs comparably to a pure VQC with amplitude embedding. This provides a more

honest baseline for evaluating quantum classifiers and serves as an important methodological contribution for the field. Unlike many optimistic reports on hybrid models, our work provides a more critical evaluation. It highlights a common pitfall in the field, where performance gains are sometimes incorrectly attributed to the quantum component. The proposed AE-VQC architecture provides a template for future research on DQCs that allows for a more transparent assessment of quantum advantage.

6.4. Comparative Analysis of Parameter Efficiency

In Kölle et al. [74], we empirically test a foundational claim of QML—that VQCs can achieve comparable performance to classical NNs while using significantly fewer trainable parameters. While this idea is a central motivation for the field, there is a lack of systematic, head-to-head comparisons that control for model size and evaluate performance across different learning paradigms. This work addresses the need to accurately test this claim in a controlled setting. We developed a unified experimental process to evaluate NNs and VQCs on both SL (classification on Iris, Wine, and WDBC datasets) and RL (QL on a custom Frozen Lake environment). For each task, a grid search is performed over different model sizes (number of layers and nodes/qubits) for both NNs and VQCs. From these results, pairs of NN and VQC models with comparable final performance are selected for a detailed analysis of their parameter counts and training times. To make the comparison more realistic, selected VQC training runs are also executed on real IBMQ hardware to estimate actual training durations.

Table 6.2.: Mean training times and parameter counts for performance-matched NNs and VQCs across SL and RL tasks. The results quantify the trade-off between the superior parameter efficiency of VQCs and their longer training times, especially on real quantum hardware. Table adapted from Kölle et al. [74].

Task	Model		Training Time (s)					
	NN	VQC	NN	VQC _{SIM}	VQC _{QPU}			
SL: Iris	NN-75	VQC-28	1.8 ± 0.0	92.6 ± 0.2	1806.1 ± 4.1			
SL: Wine	NN-105	VQC-40	1.7 ± 0.0	313.5 ± 1.3	2437.4 ± 11.6			
SL: WDBC	NN-101	VQC-63	5.4 ± 0.0	2482.9 ± 12.7	7732.5 ± 37.2			
RL	NN-112	VQC-41	95.0 ± 30.7	2511.4 ± 219.3	39330.9 ± 3458.8			

Our results, summarized in table 6.2, provide strong empirical evidence supporting the claim of quantum parameter efficiency. The study finds that, for a comparable level of performance, VQCs consistently require significantly fewer parameters than their classical NN counterparts—up to 63% fewer in some

cases. However, this parameter efficiency comes at the cost of substantially longer training times, especially when estimating execution on real quantum hardware. The main contribution is a clear, quantitative benchmark that validates the parameter efficiency advantage of VQCs while also highlighting the current practical challenge of longer training durations. This work provides a comparative study that is rarely performed with this level of detail in the QML literature. It contrasts with papers that present a new quantum model but only compare it to a single, arbitrarily chosen classical baseline. By systematically searching over model sizes for both paradigms and finding performance-matched pairs, our study offers a much fairer and more insightful comparison, grounding the theoretical promise of parameter efficiency in concrete, empirical data.

7. Conclusion

The purpose of computing is insight, not numbers.

Richard Hamming, 1962 [6]

7.1. Summary

In the current NISQ era, hybrid quantum-classical approaches like VQAs offer a promising path for enhancing AI by combining classical optimization with quantum operations. However, their practical application is limited by foundational challenges—including barren plateaus, inefficient training and the complexities of manual architecture design—which restrict performance on near-term devices. In this thesis, we address these issues through a comprehensive research framework designed to mitigate the limitations of VQAs while leveraging their unique potential in advanced AI applications. To establish a robust foundation for VQAs, we first systematically addressed bottlenecks in training and architecture design. For the training process, we introduced methods to improve robustness and efficiency, such as weight re-mapping and exponential learning rate decay. We mitigated the barren plateau phenomenon by formulating it as a problem of structural redundancy, which allowed us to prune circuits into more trainable subcircuits. As a practical alternative, we also validated the effectiveness of gradient-free optimizers. To manage the data capacity limits of NISQ devices, we introduced an architecture based on classical compression techniques. Building on these training enhancements, we replaced the manual design process with automated circuit synthesis frameworks, developing RL and EA approaches that autonomously discover high-performing, parameter-efficient VQC ansätze. With these foundational methods in place, we then demonstrated their practical value in more advanced AI applications. In RL, we used VQAs to create highly parameter-efficient agents, a key advantage for building scalable multi-agent systems. We further showcased the versatility of VQAs by designing novel quantum models for generative AI and by developing techniques to make quantum anomaly detection computationally practical. To ensure a systematic evaluation of these models, we also established a method to disentangle the contributions of hybrid systems and conducted a systematic analysis of the trade-off between parameter efficiency

and training time. While this summary presents the research as a unified narrative, it is built upon a series of individual publications created over several years. This text integrates their contributions into a comprehensive framework, unifying vocabulary and highlighting the overarching trajectory of the research. As the individual papers reflect different stages of this process, readers are encouraged to consult the appended publications (appendix A) for in-depth methodologies, extended discussions and the respective sections on related work. Following this summary, the guiding research questions introduced in section 1.2 are now revisited and concise answers are provided based on the consolidated findings of this work.

Q1 How can VQAs in the NISQ era be designed to mitigate data capacity and convergence limitations?

We showed that these limitations can be mitigated through targeted interventions at both the algorithmic and architectural levels. To address data capacity, an architectural challenge, we designed a hybrid architecture using classical compression techniques like AEs for unsupervised dimensionality reduction, which enables the processing of high-dimensional data like images and audio on hardware with a limited number of qubits. To address convergence, an algorithmic challenge, we introduced techniques that stabilize the training process. By inserting weight re-mapping that constrains parameters to their physically meaningful, periodic interval, we achieved accelerated convergence and, in several cases, statistically significant improvements in final test accuracy. We also found that systematically applying training heuristics like exponential learning rate decay is important for stable learning.

Q2 How can VQAs be made more robust and resilient to challenges like barren plateaus?

We demonstrated two complementary strategies to improve resilience to barren plateaus by addressing either the optimization method or the model architecture. First, we showed that the vanishing gradient problem can be circumvented by replacing standard optimizers with gradient-free metaheuristic algorithms. In a systematic comparison, PSO and SA emerged as particularly effective. Second, we simplified the model’s complexity through structural pruning guided by the LTH. This approach produces smaller “winning ticket” subcircuits with as few as 26% of the original parameters. We showed this method can escape a barren plateau where the full model failed, improving performance on one dataset from 45% to 80% accuracy.

Q3 How can the manual process of VQC design be automated to discover more efficient, task-specific architectures?

We demonstrated that the manual design process can be successfully automated using both RL and EAs. In our RL-based QAS framework, we showed that an agent can autonomously discover low-complexity circuits that outperform manually designed, hardware-efficient ansätze on classification tasks, as well as synthesize circuits for specific target states. Concurrently, we found

that EAs provide a powerful alternative, where a fine-grained, gate-based evolution strategy is empirically shown to be more effective than rigid, layer-based approaches at discovering compact, high-performing agents. Both paradigms systematically find efficient, task-specific architectures that rival or surpass human-designed heuristics.

Q4 How can VQCs be effectively integrated into QRL agents to address challenges in stability, optimization and scalability?

We demonstrated a successful integration of VQCs into QRL agents by introducing specific solutions for the challenges of stability, optimization and scalability. For stability, we developed a hybrid agent architecture with a classical post-processing layer that significantly improves learning performance over pure quantum approaches. With respect to optimization, we validated gradient-free metaheuristic algorithms as a robust replacement for gradient-based methods in complex control tasks. Concerning scalability, we leveraged the inherent parameter efficiency of VQCs; in a multi-agent setting, our VQC-based agents achieved performance comparable to larger classical NNs while in total using 97.88% fewer parameters.

Q5 How can quantum models be designed and adapted to address advanced generative and discriminative ML tasks?

For generative tasks, we designed parameter-efficient replacements for core components in state-of-the-art architectures. This is demonstrated through the development of a quantum-native U-Net (QU-Net) and Q-Dense network for diffusion models, which outperform classical counterparts of a similar size and were successfully demonstrated on real IBMQ hardware. For discriminative tasks, we adapted quantum models to overcome their computational scaling bottlenecks. For instance, we applied approximation techniques like random measurements and variable subsampling to quantum kernel methods for anomaly detection, reducing their time complexity from quadratic to near-linear and achieving speedups of up to 95%.

Q6 What is the empirical trade-off between the parameter efficiency of VQCs and their training time compared to classical NNs?

In our systematic, head-to-head comparison, we found the empirical trade-off to be clear: VQCs achieved a desired performance level with significantly fewer parameters than classical NNs—up to 63% fewer in controlled comparisons—thus confirming their superior parameter efficiency. However, we found that this efficiency comes at the cost of substantially longer training times. On classical simulators, VQCs are orders of magnitude slower to train and this time penalty is amplified dramatically when estimating execution on current-generation quantum hardware, identifying training duration as the primary practical bottleneck.

7.2. Outlook

The research presented in this thesis provides a comprehensive framework for designing, training and applying VQAs more effectively, moving the field from heuristic-driven exploration toward a systematic and automated methodology. While the findings address important contemporary challenges, they also highlight the significant obstacles that remain and the opportunities for future investigation. It is important to frame these future directions within the current scientific context. Unlike certain quantum algorithms with mathematically proven exponential speedups, the advantage of VQAs currently remains empirical and problem-dependent. This thesis confirmed that while VQAs offer significant parameter efficiency, this benefit is currently negated by substantial training times, rendering them impractical for most industrial applications today. Furthermore, while the techniques presented here mitigate the barren plateau phenomenon, they do not eliminate it; it remains a persistent challenge to scaling. These realities suggest that the path forward for QAI will not be a general replacement of classical methods, but rather a careful search for specific domains where quantum approaches offer a genuine advantage. Given these constraints, future research should pursue several directions of improvement. One of the most compelling directions lies in the co-design of quantum software and hardware. The automated circuit synthesis frameworks developed in this thesis could be made hardware-aware by integrating detailed noise models and qubit connectivity maps of specific devices directly into their optimization objectives. This would allow an RL agent or EA to find the optimal circuit for a particular piece of hardware, helping to bridge the simulation-to-reality gap. Algorithmically, the success of the LTH suggests a new line of inquiry: developing methods that construct sparse-by-design architectures from the outset, rather than relying on post-training pruning, to directly address the training time bottleneck. Ultimately, the contributions in this thesis integrate into a unified workflow that combines automated design, robust training and architectural pruning. This approach provides a practical method for developing effective VQAs today, while also enabling a systematic search for the specialized domains where they offer a genuine benefit.

Acronyms

A2C	Advantage Actor-Critic	11
ACO	Ant Colony Optimization	39
Adam	Adaptive Moment Estimation	8
AE	Auto-Encoder	48
AI	Artificial Intelligence	1
AUC	Area Under The Curve	9
CIFAR	Canadian Institute For Advanced Research	45
CNOT	Controlled-NOT	15
DL	Deep Learning	1
DQC	Dressed Quantum Circuit	38
DQN	Deep Q-Network	10
EA	Evolutionary Algorithm	3
FID	Frechet Inception Distance	45
GA	Genetic Algorithm	11
GELU	Gaussian Error Linear Unit	8
GMM	Gaussian Mixture Model	49
HS	Harmony Search	39
LTH	Lottery Ticket Hypothesis	22
MAE	Mean Absolute Error	9
MARL	Multi-Agent Reinforcement Learning	40
MDP	Markov Decision Process	9
ML	Machine Learning	1
MNIST	Modified National Institute Of Standards And Technology	32

MSE Mean Squared Error	9
NISQ Noisy Intermediate-Scale Quantum	2
NN Artificial Neural Network	2
OC-SVM One-Class Support Vector Machine	46
PCA Principal Component Analysis	49
PPO Proximal Policy Optimization	11
PQC Parametrized Quantum Circuit	18
PSNR Peak Signal-to-Noise Ratio	45
PSO Particle Swarm Optimization	39
QAI Quantum Artificial Intelligence	3
QAOA Quantum Approximate Optimization Algorithm	17
QAS Quantum Architecture Search	32
QC Quantum Computing	1
QL Q-Learning	10
QMARL Quantum Multi-Agent Reinforcement Learning	33
QML Quantum Machine Learning	2
QPPO Quantum Proximal Policy Optimization	24
QRL Quantum Reinforcement Learning	3
QU-Net Quantum U-Net	45
Qubit Quantum Bit	1
ReLU Rectified Linear Unit	8
RL Reinforcement Learning	3
RMSE Root Mean Squared Error	9
RMSPProp Root Mean Square Propagation	8
ROC Receiver Operating Characteristic	9
SA Simulated Annealing	11
SEL Strongly Entangling Layer	33
SGD Stochastic Gradient Descent	8
SL Supervised Learning	8

SSIM	Structural Similarity Index Measure	45
SVM	Support Vector Machine	46
TD	Temporal-Difference	9
TS	Tabu Search	11
VQA	Variational Quantum Algorithm	2
VQC	Variational Quantum Circuit	2

List of Figures

2.1. Bloch sphere representation	14
2.2. Example quantum circuit	16
2.3. VQA structure	18
2.4. VQC structure and data flow	19
3.1. Overview of the VQC training process with weight re-mapping .	22
3.2. Training results for weight re-mapping	23
3.3. Training Heuristics Overview	25
3.4. Performance of training heuristics in CartPole and FrozenLake .	26
3.5. Performance of pruned VQC on the Wine dataset	27
4.1. Agent performance scaling with qubits and circuit depth	31
4.2. RL-QAS Framework Overview	32
4.3. Comparison of evolutionary architectural approaches	34
4.4. Histogram of fitness scores for mutation strategies	35
5.1. Hybrid VQC architectures	38
5.2. HA2C performance comparison	39
5.3. Metaheuristic algorithm performance comparison	40
5.4. Multi-agent VQC performance	41
6.1. Iterative denoising process of the Q-Dense model	44
6.2. QU-Net architecture and performance	45
6.3. Training and testing time for anomaly detection methods	46
6.4. Performance metrics for anomaly detection methods	47
6.5. Hybrid AE-VQC architecture	48

Bibliography

- [1] Emile H. L. Aarts and Jan K. Lenstra. *Local Search in Combinatorial Optimization*. Princeton University Press, Sept. 2003. DOI: [10.2307/j.ctv346t9c](https://doi.org/10.2307/j.ctv346t9c). URL: <http://www.jstor.org/stable/j.ctv346t9c>.
- [2] Amira Abbas, David Sutter, Christa Zoufal, Aurelien Lucchi, Alessio Figalli, and Stefan Woerner. “The power of quantum neural networks”. In: *Nature Computational Science* 1.6 (June 2021), pp. 403–409. ISSN: 2662-8457. DOI: [10.1038/s43588-021-00084-1](https://doi.org/10.1038/s43588-021-00084-1). URL: <https://doi.org/10.1038/s43588-021-00084-1>.
- [3] Muhammad AbuGhanem. “IBM quantum computers: evolution, performance, and future directions”. In: *The Journal of Supercomputing* 81.5 (Apr. 2025), p. 687. ISSN: 1573-0484. DOI: [10.1007/s11227-025-07047-7](https://doi.org/10.1007/s11227-025-07047-7). URL: <https://doi.org/10.1007/s11227-025-07047-7>.
- [4] Philipp Altmann, Leo Sünkel, Jonas Stein, Tobias Müller, Christoph Roch, and Claudia Linnhoff-Popien. “SEQUENT: Towards Traceable Quantum Machine Learning Using Sequential Quantum Enhanced Training”. In: *Proceedings of the 15th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART*. INSTICC. SciTePress, 2023, pp. 744–751. ISBN: 978-989-758-623-1. DOI: [10.5220/0011772400003393](https://doi.org/10.5220/0011772400003393).
- [5] Dario Amodei and Danny Hernandez. *AI and compute*. <https://openai.com/index/ai-and-compute/>. Accessed on 08/30/2025. May 2018.
- [6] H. M. Antia. *Numerical Methods for Scientists and Engineers*. Hindustan Book Agency, 2012. DOI: [10.1007/978-93-86279-52-1](https://doi.org/10.1007/978-93-86279-52-1). URL: <http://dx.doi.org/10.1007/978-93-86279-52-1>.
- [7] *Artificial Intelligence Market Size, Share & Trends Analysis Report by Solution, by Technology (Deep Learning, Machine Learning, NLP, Machine Vision, Generative AI), by Function, by End-Use, by Region, and Segment Forecasts, 2025–2030*. Tech. rep. GVR-1-68038-955-5. Accessed on 09/07/2025. Grand View Research, 2025. URL: <https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-market>.
- [8] Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J. Fleet. *Synthetic Data from Diffusion Models Improves ImageNet Classification*. 2023. arXiv: [2304.08466](https://arxiv.org/abs/2304.08466) [cs.CV]. URL: <https://arxiv.org/abs/2304.08466>.

- [9] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. “Elementary gates for quantum computation”. In: *Phys. Rev. A* 52 (5 Nov. 1995), pp. 3457–3467. DOI: [10.1103/PhysRevA.52.3457](https://doi.org/10.1103/PhysRevA.52.3457). URL: <https://link.aps.org/doi/10.1103/PhysRevA.52.3457>.
- [10] J. S. Bell. “On the Einstein Podolsky Rosen paradox”. In: *Physics Physique Fizika* 1 (3 Nov. 1964), pp. 195–200. DOI: [10.1103/PhysicsPhysiqueFizika.1.195](https://doi.org/10.1103/PhysicsPhysiqueFizika.1.195). URL: <https://link.aps.org/doi/10.1103/PhysicsPhysiqueFizika.1.195>.
- [11] Richard Bellman. “Dynamic programming and stochastic control processes”. In: *Information and Control* 1.3 (1958), pp. 228–239. ISSN: 0019-9958. DOI: [10.1016/S0019-9958\(58\)80003-0](https://doi.org/10.1016/S0019-9958(58)80003-0). URL: <https://www.sciencedirect.com/science/article/pii/S0019995858800030>.
- [12] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. “Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels”. In: *Phys. Rev. Lett.* 70 (13 Mar. 1993), pp. 1895–1899. DOI: [10.1103/PhysRevLett.70.1895](https://doi.org/10.1103/PhysRevLett.70.1895). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.70.1895>.
- [13] Kishor Bharti, Alba Cervera-Liarta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik. “Noisy intermediate-scale quantum algorithms”. In: *Rev. Mod. Phys.* 94 (1 Feb. 2022), p. 015004. DOI: [10.1103/RevModPhys.94.015004](https://doi.org/10.1103/RevModPhys.94.015004). URL: <https://link.aps.org/doi/10.1103/RevModPhys.94.015004>.
- [14] Sounak Bhowmik, Talita Perciano, and Himanshu Thapliyal. “Quantum Transfer Learning to Boost Dementia Detection”. In: *Proceedings of the Great Lakes Symposium on VLSI 2025*. GLSVLSI ’25. Association for Computing Machinery, 2025, pp. 849–853. ISBN: 9798400714962. DOI: [10.1145/3716368.3735290](https://doi.org/10.1145/3716368.3735290). URL: <https://doi.org/10.1145/3716368.3735290>.
- [15] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. “Quantum machine learning”. In: *Nature* 549.7671 (Sept. 2017), pp. 195–202. ISSN: 1476-4687. DOI: [10.1038/nature23474](https://doi.org/10.1038/nature23474). URL: <https://doi.org/10.1038/nature23474>.
- [16] Christopher M. Bishop. *Pattern recognition and machine learning*. New York : Springer, [2006] ©2006, 2006. DOI: [10.1007/978-0-387-45528-0](https://doi.org/10.1007/978-0-387-45528-0). URL: <https://search.library.wisc.edu/catalog/9910032530902121>.

-
- [17] Joseph Bowles, Shahnawaz Ahmed, and Maria Schuld. *Better than classical? The subtle art of benchmarking quantum machine learning models*. 2024. arXiv: [2403.07059](https://arxiv.org/abs/2403.07059) [quant-ph]. URL: <https://arxiv.org/abs/2403.07059>.
- [18] Benjamin Brewster. “Portfolio: Theory and Practice”. In: *The Yale Literary Magazine* 47.5 (Feb. 1882), p. 202.
- [19] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. “Language models are few-shot learners”. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS ’20. Vancouver, BC, Canada: Curran Associates Inc., 2020. ISBN: 9781713829546. DOI: [10.5555/3495724.3495883](https://doi.org/10.5555/3495724.3495883).
- [20] Bruce Buchanan and Edward Shortliffe. *Rule-based Expert System – The MYCIN Experiments of the Stanford Heuristic Programming Project*. Vol. 3. 4. Elsevier BV, Jan. 1986, p. 304. DOI: [10.1016/s0736-5853\(86\)80073-9](https://doi.org/10.1016/s0736-5853(86)80073-9). URL: [http://dx.doi.org/10.1016/s0736-5853\(86\)80073-9](http://dx.doi.org/10.1016/s0736-5853(86)80073-9).
- [21] Lucian Busoniu, Robert Babuska, and Bart De Schutter. “A Comprehensive Survey of Multiagent Reinforcement Learning”. In: *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 38 (Apr. 2008), pp. 156–172. DOI: [10.1109/TSMCC.2007.913919](https://doi.org/10.1109/TSMCC.2007.913919).
- [22] Zhenyu Cai, Ryan Babbush, Simon C. Benjamin, Suguru Endo, William J. Huggins, Ying Li, Jarrod R. McClean, and Thomas E. O’Brien. “Quantum error mitigation”. In: *Rev. Mod. Phys.* 95 (4 Dec. 2023), p. 045005. DOI: [10.1103/RevModPhys.95.045005](https://link.aps.org/doi/10.1103/RevModPhys.95.045005). URL: <https://link.aps.org/doi/10.1103/RevModPhys.95.045005>.
- [23] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3.9 (Sept. 2021), pp. 625–644. ISSN: 2522-5820. DOI: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9). URL: <https://doi.org/10.1038/s42254-021-00348-9>.

- [24] T. Chai and R. R. Draxler. “Root mean square error (RMSE) or mean absolute error (MAE)? – Arguments against avoiding RMSE in the literature”. In: *Geoscientific Model Development* 7.3 (2014), pp. 1247–1250. DOI: [10.5194/gmd-7-1247-2014](https://doi.org/10.5194/gmd-7-1247-2014). URL: <https://gmd.copernicus.org/articles/7/1247/2014/>.
- [25] Samuel Yen-Chi Chen. “An Introduction to Quantum Reinforcement Learning (QRL)”. In: *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*. 2024, pp. 1139–1144. DOI: [10.1109/ICTC62082.2024.10827651](https://doi.org/10.1109/ICTC62082.2024.10827651).
- [26] Samuel Yen-Chi Chen. “Differentiable Quantum Architecture Search in Asynchronous Quantum Reinforcement Learning”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Vol. 01. 2024, pp. 1516–1524. DOI: [10.1109/QCE60285.2024.00178](https://doi.org/10.1109/QCE60285.2024.00178).
- [27] Samuel Yen-Chi Chen. “Quantum Reinforcement Learning for Quantum Architecture Search”. In: *Proceedings of the 2023 International Workshop on Quantum Classical Cooperative*. QCCC ’23. Orlando, FL, USA: Association for Computing Machinery, 2023, pp. 17–20. ISBN: 9798400701627. DOI: [10.1145/3588983.3596692](https://doi.org/10.1145/3588983.3596692). URL: <https://doi.org/10.1145/3588983.3596692>.
- [28] Samuel Yen-Chi Chen, Chih-Min Huang, Chia-Wei Hsing, Hsi-Sheng Goan, and Ying-Jer Kao. “Variational quantum reinforcement learning via evolutionary optimization”. In: *Machine Learning: Science and Technology* 3.1 (Feb. 2022), p. 015025. DOI: [10.1088/2632-2153/ac4559](https://doi.org/10.1088/2632-2153/ac4559). URL: <https://dx.doi.org/10.1088/2632-2153/ac4559>.
- [29] Samuel Yen-Chi Chen, Chao-Han Huck Yang, Jun Qi, Pin-Yu Chen, Xiaoli Ma, and Hsi-Sheng Goan. “Variational Quantum Circuits for Deep Reinforcement Learning”. In: *IEEE Access* 8 (2020), pp. 141007–141024. DOI: [10.1109/ACCESS.2020.3010470](https://doi.org/10.1109/ACCESS.2020.3010470).
- [30] A. D. Córcoles, Maika Takita, Ken Inoue, Scott Lekuch, Zlatko K. Mineev, Jerry M. Chow, and Jay M. Gambetta. “Exploiting Dynamic Quantum Circuits in a Quantum Algorithm with Superconducting Qubits”. In: *Phys. Rev. Lett.* 127 (10 Aug. 2021), p. 100501. DOI: [10.1103/PhysRevLett.127.100501](https://doi.org/10.1103/PhysRevLett.127.100501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.100501>.
- [31] Daniel Crevier. *AI: the tumultuous history of the search for artificial intelligence*. USA: Basic Books, Inc., 1993. ISBN: 0465029973. DOI: [10.5555/151188](https://doi.org/10.5555/151188).
- [32] G. Cybenko. “Approximation by superpositions of a sigmoidal function”. In: *Mathematics of Control, Signals and Systems* 2.4 (Dec. 1989), pp. 303–314. ISSN: 1435-568X. DOI: [10.1007/BF02551274](https://doi.org/10.1007/BF02551274). URL: <https://doi.org/10.1007/BF02551274>.

- [33] Xin Dai, Tzu-Chieh Wei, Shinjae Yoo, and Samuel Yen-Chi Chen. “Quantum Machine Learning Architecture Search via Deep Reinforcement Learning”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Vol. 01. 2024, pp. 1525–1534. DOI: [10.1109/QCE60285.2024.00179](https://doi.org/10.1109/QCE60285.2024.00179).
- [34] Andrew J. Daley, Immanuel Bloch, Christian Kokail, Stuart Flannigan, Natalie Pearson, Matthias Troyer, and Peter Zoller. “Practical quantum advantage in quantum simulation”. en. In: *Nature* 607.7920 (July 2022), pp. 667–676. DOI: [10.1038/s41586-022-04940-6](https://doi.org/10.1038/s41586-022-04940-6). URL: <http://dx.doi.org/10.1038/s41586-022-04940-6>.
- [35] David Deutsch. “Quantum theory, the Church–Turing principle and the universal quantum computer”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 400.1818 (1985), pp. 97–117. ISSN: 0080-4630. DOI: [10.1098/rspa.1985.0070](https://doi.org/10.1098/rspa.1985.0070).
- [36] Petros Drineas and Michael W. Mahoney. “On the Nyström Method for Approximating a Gram Matrix for Improved Kernel-Based Learning”. In: *J. Mach. Learn. Res.* 6 (Dec. 2005), pp. 2153–2175. ISSN: 1532-4435. DOI: [10.5555/1046920.1194916](https://doi.org/10.5555/1046920.1194916).
- [37] Vedran Dunjko and Hans J Briegel. “Machine learning & artificial intelligence in the quantum domain: a review of recent progress”. In: *Reports on Progress in Physics* 81.7 (June 2018), p. 074001. DOI: [10.1088/1361-6633/aab406](https://doi.org/10.1088/1361-6633/aab406). URL: <http://dx.doi.org/10.1088/1361-6633/aab406>.
- [38] Bryan Eastin and Emanuel Knill. “Restrictions on Transversal Encoded Quantum Gate Sets”. In: *Phys. Rev. Lett.* 102 (11 Mar. 2009), p. 110502. DOI: [10.1103/PhysRevLett.102.110502](https://doi.org/10.1103/PhysRevLett.102.110502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.110502>.
- [39] Jens Eisert, Martin Wilkens, and Maciej Lewenstein. “Quantum Games and Quantum Strategies”. In: *Phys. Rev. Lett.* 83 (15 Oct. 1999), pp. 3077–3080. DOI: [10.1103/PhysRevLett.83.3077](https://doi.org/10.1103/PhysRevLett.83.3077). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.83.3077>.
- [40] Artur K. Ekert. “Quantum cryptography based on Bell’s theorem”. In: *Phys. Rev. Lett.* 67 (6 Aug. 1991), pp. 661–663. DOI: [10.1103/PhysRevLett.67.661](https://doi.org/10.1103/PhysRevLett.67.661). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.67.661>.
- [41] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. “Dermatologist-level classification of skin cancer with deep neural networks”. en. In: *Nature* 542.7639 (Jan. 2017), pp. 115–118. DOI: [10.1038/nature21056](https://doi.org/10.1038/nature21056). URL: <http://dx.doi.org/10.1038/nature21056>.

- [42] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. *A Quantum Approximate Optimization Algorithm*. 2014. arXiv: [1411 . 4028 \[quant-ph\]](https://arxiv.org/abs/1411.4028). URL: <https://arxiv.org/abs/1411.4028>.
- [43] Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Francisco J. R. Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, David Silver, Demis Hassabis, and Pushmeet Kohli. “Discovering faster matrix multiplication algorithms with reinforcement learning”. In: *Nature* 610.7930 (Oct. 2022), pp. 47–53. ISSN: 1476-4687. DOI: [10 . 1038 / s41586 - 022 - 05172 - 4](https://doi.org/10.1038/s41586-022-05172-4). URL: <https://doi.org/10.1038/s41586-022-05172-4>.
- [44] Richard P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6 (June 1982), pp. 467–488. ISSN: 1572-9575. DOI: [10.1007/BF02650179](https://doi.org/10.1007/BF02650179). URL: <https://doi.org/10.1007/BF02650179>.
- [45] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. “Surface codes: Towards practical large-scale quantum computation”. In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: [10.1103/PhysRevA.86.032324](https://link.aps.org/doi/10.1103/PhysRevA.86.032324). URL: <https://link.aps.org/doi/10.1103/PhysRevA.86.032324>.
- [46] T. Gabor, M. Zorn, and Linnhoff-Popien C. “The applicability of reinforcement learning for the automatic generation of state preparation circuits”. In: *GECCO ’22: Proceedings of the Genetic and Evolutionary Computation Conference Companion* (2022), pp. 2196–2204. DOI: [10.1145/3520304.3534039](https://doi.org/10.1145/3520304.3534039).
- [47] Alessandro Giovagnoli, Yunpu Ma, and Volker Tresp. *QNEAT: Natural Evolution of Variational Quantum Circuit Architecture*. 2023. arXiv: [2304.06981 \[quant-ph\]](https://arxiv.org/abs/2304.06981). URL: <https://arxiv.org/abs/2304.06981>.
- [48] James Gleick and Freeman J. Dyson. “Genius: The Life and Science of Richard Feynman”. en. In: *Physics Today* 45.11 (Nov. 1992), pp. 87–87. DOI: [10.1063/1.2809877](https://dx.doi.org/10.1063/1.2809877). URL: <http://dx.doi.org/10.1063/1.2809877>.
- [49] Fred Glover. “Tabu Search—Part I”. en. In: *ORSA Journal on Computing* 1.3 (Aug. 1989), pp. 190–206. DOI: [10.1287/ijoc.1.3.190](https://dx.doi.org/10.1287/ijoc.1.3.190). URL: <http://dx.doi.org/10.1287/ijoc.1.3.190>.
- [50] David E. Goldberg and John H. Holland. “Genetic Algorithms and Machine Learning”. In: *Machine Learning* 3.2 (Oct. 1988), pp. 95–99. ISSN: 1573-0565. DOI: [10.1023/A:1022602019183](https://doi.org/10.1023/A:1022602019183). URL: <https://doi.org/10.1023/A:1022602019183>.
- [51] Daniel Gottesman. *Stabilizer Codes and Quantum Error Correction*. 1997. arXiv: [quant-ph/9705052 \[quant-ph\]](https://arxiv.org/abs/quant-ph/9705052). URL: <https://arxiv.org/abs/quant-ph/9705052>.

- [52] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. “An adaptive variational algorithm for exact molecular simulations on a quantum computer”. en. In: *Nature Communications* 10.1 (July 2019). DOI: [10.1038/s41467-019-10988-2](https://doi.org/10.1038/s41467-019-10988-2). URL: <http://dx.doi.org/10.1038/s41467-019-10988-2>.
- [53] Lov 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. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 212–219. ISBN: 0897917855. DOI: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866). URL: <https://doi.org/10.1145/237814.237866>.
- [54] *Handbook of Metaheuristics*. en. Springer US, 2010. DOI: [10.1007/978-1-4419-1665-5](https://doi.org/10.1007/978-1-4419-1665-5). URL: <http://dx.doi.org/10.1007/978-1-4419-1665-5>.
- [55] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. “Supervised learning with quantum-enhanced feature spaces”. In: *Nature* 567.7747 (Mar. 2019), pp. 209–212. ISSN: 1476-4687. DOI: [10.1038/s41586-019-0980-2](https://doi.org/10.1038/s41586-019-0980-2). URL: <https://doi.org/10.1038/s41586-019-0980-2>.
- [56] J. Helsen, I. Roth, E. Onorati, A.H. Werner, and J. Eisert. “General Framework for Randomized Benchmarking”. In: *PRX Quantum* 3.2 (June 2022). ISSN: 2691-3399. DOI: [10.1103/prxquantum.3.020357](https://doi.org/10.1103/prxquantum.3.020357). URL: <http://dx.doi.org/10.1103/PRXQuantum.3.020357>.
- [57] Maxwell Henderson, Samriddhi Shakya, Shashindra Pradhan, and Tristan Cook. “Quanvolutional neural networks: powering image recognition with quantum circuits”. In: *Quantum Machine Intelligence* 2.1 (Feb. 2020), p. 2. ISSN: 2524-4914. DOI: [10.1007/s42484-020-00012-y](https://doi.org/10.1007/s42484-020-00012-y). URL: <https://doi.org/10.1007/s42484-020-00012-y>.
- [58] Dan Hendrycks and Kevin Gimpel. *Gaussian Error Linear Units (GELUs)*. 2023. arXiv: [1606.08415](https://arxiv.org/abs/1606.08415) [cs.LG]. URL: <https://arxiv.org/abs/1606.08415>.
- [59] Jonathan Ho, Ajay Jain, and Pieter Abbeel. “Denoising diffusion probabilistic models”. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS '20. Vancouver, BC, Canada: Curran Associates Inc., 2020. ISBN: 9781713829546. DOI: [10.5555/3495724.3496298](https://doi.org/10.5555/3495724.3496298).
- [60] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. “Multilayer feedforward networks are universal approximators”. In: *Neural Networks* 2.5 (1989), pp. 359–366. ISSN: 0893-6080. DOI: [10.1016/0893-6080\(89\)90020-8](https://www.sciencedirect.com/science/article/pii/0893608089900208). URL: <https://www.sciencedirect.com/science/article/pii/0893608089900208>.

- [61] Dominic Horsman, Austin G Fowler, Simon Devitt, and Rodney Van Meter. “Surface code quantum computing by lattice surgery”. In: *New Journal of Physics* 14.12 (Dec. 2012), p. 123011. DOI: [10.1088/1367-2630/14/12/123011](https://doi.org/10.1088/1367-2630/14/12/123011). URL: <https://dx.doi.org/10.1088/1367-2630/14/12/123011>.
- [62] Ronald A. Howard. *Dynamic programming and Markov processes*. Oxford, England: John Wiley, 1960, pp. viii, 136–viii, 136.
- [63] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Predicting many properties of a quantum system from very few measurements”. In: *Nature Physics* 16.10 (Oct. 2020), pp. 1050–1057. ISSN: 1745-2481. DOI: [10.1038/s41567-020-0932-7](https://doi.org/10.1038/s41567-020-0932-7). URL: <https://doi.org/10.1038/s41567-020-0932-7>.
- [64] Sofiene Jerbi, Lea M. Trenkwalder, Hendrik Poulsen Nautrup, Hans J. Briegel, and Vedran Dunjko. “Quantum Enhancements for Deep Reinforcement Learning in Large Spaces”. In: *PRX Quantum* 2 (1 Feb. 2021), p. 010328. DOI: [10.1103/PRXQuantum.2.010328](https://link.aps.org/doi/10.1103/PRXQuantum.2.010328). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.010328>.
- [65] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. “Highly accurate protein structure prediction with AlphaFold”. In: *Nature* 596.7873 (Aug. 2021), pp. 583–589. ISSN: 1476-4687. DOI: [10.1038/s41586-021-03819-2](https://doi.org/10.1038/s41586-021-03819-2). URL: <https://doi.org/10.1038/s41586-021-03819-2>.
- [66] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. In: *Nature* 549.7671 (Sept. 2017), pp. 242–246. ISSN: 1476-4687. DOI: [10.1038/nature23879](https://doi.org/10.1038/nature23879). URL: <https://doi.org/10.1038/nature23879>.
- [67] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. *Scaling Laws for Neural Language Models*. 2020. arXiv: [2001.08361](https://arxiv.org/abs/2001.08361) [cs.LG]. URL: <https://arxiv.org/abs/2001.08361>.
- [68] Dohun Kim and Seokhyeong Kang. “Quantum Denoising Diffusion Probabilistic Models for Image Generation”. In: *Korean Conference on Semiconductors*. Vol. 30. Feb. 2023.

-
- [69] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG]. URL: <https://arxiv.org/abs/1412.6980>.
- [70] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. “Optimization by Simulated Annealing”. In: *Science* 220.4598 (1983), pp. 671–680. DOI: [10.1126/science.220.4598.671](https://doi.org/10.1126/science.220.4598.671). eprint: <https://www.science.org/doi/pdf/10.1126/science.220.4598.671>. URL: <https://www.science.org/doi/abs/10.1126/science.220.4598.671>.
- [71] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver. “Superconducting Qubits: Current State of Play”. en. In: *Annual Review of Condensed Matter Physics* 11.1 (Mar. 2020), pp. 369–395. DOI: [10.1146/annurev-conmatphys-031119-050605](https://doi.org/10.1146/annurev-conmatphys-031119-050605). URL: <http://dx.doi.org/10.1146/annurev-conmatphys-031119-050605>.
- [72] Michael Kölle, Afrae Ahouzi, Pascal Debus, Robert Müller, Daniëlle Schuman, and Claudia Linnhoff-Popien. “Towards Efficient Quantum Anomaly Detection: One-Class SVMs Using Variable Subsampling and Randomized Measurements”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*. INSTICC. SciTePress, 2024, pp. 324–335. ISBN: 978-989-758-680-4. DOI: [10.5220/0012381200003636](https://doi.org/10.5220/0012381200003636). URL: <https://www.scitepress.org/Papers/2024/123812/>.
- [73] Michael Kölle, Tom Bintener, Maximilian Zorn, Gerhard Stenzel, Leo Sünkel, Thomas Gabor, and Claudia Linnhoff-Popien. “Evaluating Mutation Techniques in Genetic Algorithm-Based Quantum Circuit Synthesis”. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. GECCO '25. NH Malaga Hotel, Malaga, Spain: Association for Computing Machinery, July 2025, pp. 907–915. ISBN: 9798400714658. DOI: [10.1145/3712256.3726402](https://doi.org/10.1145/3712256.3726402). URL: <https://doi.org/10.1145/3712256.3726402>.
- [74] Michael Kölle, Alexander Feist, Jonas Stein, Sebastian Wölckert, and Claudia Linnhoff-Popien. “Evaluating Parameter-Based Training Performance of Neural Networks and Variational Quantum Circuits”. In: *Computational Science – ICCS 2025*. Ed. by Michael H. Lees, Wentong Cai, Siew Ann Cheong, Yi Su, David Abramson, Jack J. Dongarra, and Peter M. A. Sloot. Cham: Springer Nature Switzerland, 2025, pp. 275–282. DOI: [10.1007/978-3-031-97635-3_33](https://doi.org/10.1007/978-3-031-97635-3_33).
- [75] Michael Kölle, Alessandro Giovagnoli, Jonas Stein, Maximilian Mansky, Julian Hager, and Claudia Linnhoff-Popien. “Improving Convergence for Quantum Variational Classifiers Using Weight Re-Mapping”. In: *Proceedings of the 15th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*. INSTICC. SciTePress, 2023,

- pp. 251–258. ISBN: 978-989-758-623-1. DOI: [10.5220/0011696300003393](https://doi.org/10.5220/0011696300003393). URL: <https://www.scitepress.org/Papers/2023/116963/>.
- [76] Michael Kölle, Alessandro Giovagnoli, Jonas Stein, Maximilian Balthasar Mansky, Julian Hager, Tobias Rohe, Robert Müller, and Claudia Linnhoff-Popien. “Weight Re-mapping for Variational Quantum Algorithms”. In: *Agents and Artificial Intelligence: 15th International Conference, ICAART 2023, Lisbon, Portugal, February 22–24, 2023, Revised Selected Papers*. Lisbon, Portugal: Springer-Verlag, 2024, pp. 286–309. ISBN: 978-3-031-55325-7. DOI: [10.1007/978-3-031-55326-4_14](https://doi.org/10.1007/978-3-031-55326-4_14). URL: https://doi.org/10.1007/978-3-031-55326-4_14.
- [77] Michael Kölle, Mohamad Hgog, Fabian Ritz, Philipp Altmann, Maximilian Zorn, Jonas Stein, and Claudia Linnhoff-Popien. “Quantum Advantage Actor-Critic for Reinforcement Learning”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART*. INSTICC. SciTePress, 2024, pp. 297–304. ISBN: 978-989-758-680-4. DOI: [10.5220/0012383900003636](https://doi.org/10.5220/0012383900003636). URL: <https://www.scitepress.org/Papers/2024/123839/>.
- [78] Michael Kölle, Leonhard Klingert, Julian Schönberger, Philipp Altmann, Tobias Rohe, and Claudia Linnhoff-Popien. “Investigating the Lottery Ticket Hypothesis for Variational Quantum Circuits”. In: *Proceedings of the 18th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART*. INSTICC. SciTePress, Mar. 2026, pp. 2433–2440. ISBN: 978-989-758-796-2. DOI: [10.5220/0014247600004052](https://doi.org/10.5220/0014247600004052). URL: <https://www.scitepress.org/Link.aspx?doi=10.5220/0014247600004052>.
- [79] Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, Julian Hager, Sebastian Zielinski, and Claudia Linnhoff-Popien. “Quantum Transfer Learning: The Impact of Classical Preprocessing”. In: ed. by Ana Paula Rocha, Luc Steels, and Jaap van den Herik. Cham: Springer Nature Switzerland, Apr. 2025, pp. 368–403. ISBN: 978-3-031-87327-0. DOI: [10.1007/978-3-031-87327-0_18](https://doi.org/10.1007/978-3-031-87327-0_18).
- [80] Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, and Claudia Linnhoff-Popien. “Disentangling Quantum and Classical Contributions in Hybrid Quantum Machine Learning Architectures”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART*. INSTICC. SciTePress, 2024, pp. 649–656. ISBN: 978-989-758-680-4. DOI: [10.5220/0012381600003636](https://doi.org/10.5220/0012381600003636). URL: <https://www.scitepress.org/Papers/2024/123816/>.

- [81] Michael Kölle, Simon Salfer, Tobias Rohe, Philipp Altmann, and Claudia Linnhoff-Popien. “Quantum Architecture Search for Solving Quantum Machine Learning Tasks”. In: *Proceedings of the 18th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*. INSTICC. SciTePress, Mar. 2026, pp. 1254–1263. ISBN: 978-989-758-796-2. DOI: [10.5220/0014247700004052](https://doi.org/10.5220/0014247700004052). URL: <https://www.scitepress.org/Link.aspx?doi=10.5220/0014247700004052>.
- [82] Michael Kölle, Karola Schneider, Sabrina Egger, Felix Topp, Thomy Phan, Philipp Altmann, Jonas Nüßlein, and Claudia Linnhoff-Popien. “Architectural Influence on Variational Quantum Circuits in Multi-Agent Reinforcement Learning: Evolutionary Strategies for Optimization”. In: ed. by Ana Paula Rocha, Luc Steels, and Jaap van den Herik. Cham: Springer Nature Switzerland, Apr. 2025, pp. 50–79. ISBN: 978-3-031-87327-0. DOI: [10.1007/978-3-031-87327-0_3](https://doi.org/10.1007/978-3-031-87327-0_3). eprint: [2407.20739](https://doi.org/2407.20739).
- [83] Michael Kölle, Tom Schubert, Philipp Altmann, Maximilian Zorn, Jonas Stein, and Claudia Linnhoff-Popien. “A Reinforcement Learning Environment for Directed Quantum Circuit Synthesis”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART*. INSTICC. SciTePress, 2024, pp. 83–94. ISBN: 978-989-758-680-4. DOI: [10.5220/0012383200003636](https://doi.org/10.5220/0012383200003636). URL: <https://www.scitepress.org/Papers/2024/123832/>.
- [84] Michael Kölle, Daniel Seidl, Maximilian Zorn, Philipp Altmann, Jonas Stein, and Thomas Gabor. “Optimizing Variational Quantum Circuits Using Metaheuristic Strategies in Reinforcement Learning”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Los Alamitos, CA, USA: IEEE Computer Society, Sept. 2024, pp. 323–328. DOI: [10.1109/QCE60285.2024.10300](https://doi.org/10.1109/QCE60285.2024.10300). URL: <https://www.computer.org/csdl/proceedings-article/qce/2024/413702a323/23oqmUz6Jj0>.
- [85] Michael Kölle, Gerhard Stenzel, Jonas Stein, Sebastian Zielinski, Bjorn Ommer, and Claudia Linnhoff-Popien. “Quantum Denoising Diffusion Models”. In: *2024 IEEE International Conference on Quantum Software (QSW)*. Los Alamitos, CA, USA: IEEE Computer Society, July 2024, pp. 88–98. DOI: [10.1109/QSW62656.2024.00023](https://doi.org/10.1109/QSW62656.2024.00023). URL: <https://www.computer.org/csdl/proceedings-article/qsw/2024/684700a088/1zMeeEZjPPy>.
- [86] Michael Kölle, Felix Topp, Thomy Phan, Philipp Altmann, Jonas Nüßlein, and Claudia Linnhoff-Popien. “Multi-Agent Quantum Reinforcement Learning Using Evolutionary Optimization”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART*. INSTICC. SciTePress, 2024,

- pp. 71–82. ISBN: 978-989-758-680-4. DOI: [10.5220/0012382800003636](https://doi.org/10.5220/0012382800003636). URL: <https://www.scitepress.org/Papers/2024/123828>.
- [87] Michael Kölle, Timo Witter, Tobias Rohe, Gerhard Stenzel, Philipp Altmann, and Thomas Gabor. “A Study on Optimization Techniques for Variational Quantum Circuits in Reinforcement Learning”. In: *2024 IEEE International Conference on Quantum Software (QSW)*. Los Alamitos, CA, USA: IEEE Computer Society, July 2024, pp. 157–167. DOI: [10.1109/QSW62656.2024.00031](https://doi.org/10.1109/QSW62656.2024.00031). URL: <https://www.computer.org/csdl/proceedings-article/qsw/2024/684700a157/1ZMefLwAX1m>.
- [88] Mario Krenn, Jakob S. Kottmann, Nora Tischler, and Alán Aspuru-Guzik. “Conceptual Understanding through Efficient Automated Design of Quantum Optical Experiments”. In: *Phys. Rev. X* 11 (2021). ISSN: 2160-3308. DOI: [10.1103/physrevx.11.031044](https://doi.org/10.1103/physrevx.11.031044).
- [89] Mario Krenn, Mehul Malik, Robert Fickler, Radek Lapkiewicz, and Anton Zeilinger. “Automated Search for new Quantum Experiments”. In: *Phys. Rev. Lett.* 116, 090405 (2016) 116.9 (Mar. 2015), p. 090405. DOI: [10.1103/physrevlett.116.090405](https://doi.org/10.1103/physrevlett.116.090405). arXiv: [1509.02749](https://arxiv.org/abs/1509.02749) [quant-ph].
- [90] Oleksandr Kyriienko and Einar B. Magnusson. *Unsupervised quantum machine learning for fraud detection*. 2022. arXiv: [2208.01203](https://arxiv.org/abs/2208.01203) [quant-ph]. URL: <https://arxiv.org/abs/2208.01203>.
- [91] Harashta Tatimma Larasati, Thi-Thu-Huong Le, and Howon Kim. “Trends of Quantum Computing Applications to Computer Vision”. In: *2022 International Conference on Platform Technology and Service (PlatCon)*. 2022, pp. 7–12. DOI: [10.1109/PlatCon55845.2022.9932103](https://doi.org/10.1109/PlatCon55845.2022.9932103).
- [92] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *Nature* 521.7553 (May 2015), pp. 436–444. ISSN: 1476-4687. DOI: [10.1038/nature14539](https://doi.org/10.1038/nature14539). URL: <https://doi.org/10.1038/nature14539>.
- [93] Lorenzo Leone, Salvatore F.E. Oliviero, Lukasz Cincio, and M. Cerezo. “On the practical usefulness of the Hardware Efficient Ansatz”. In: *Quantum* 8 (July 2024), p. 1395. ISSN: 2521-327X. DOI: [10.22331/q-2024-07-03-1395](https://doi.org/10.22331/q-2024-07-03-1395). URL: <http://dx.doi.org/10.22331/q-2024-07-03-1395>.
- [94] Gushu Li, Yufei Ding, and Yuan Xie. “Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices”. In: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’19. Providence, RI, USA: Association for Computing Machinery, 2019, pp. 1001–1014. ISBN: 9781450362405. DOI: [10.1145/3297858.3304023](https://doi.org/10.1145/3297858.3304023). URL: <https://doi.org/10.1145/3297858.3304023>.

-
- [95] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. “A rigorous and robust quantum speed-up in supervised machine learning”. In: *Nature Physics* 17.9 (Sept. 2021), pp. 1013–1017. ISSN: 1745-2481. DOI: [10.1038/s41567-021-01287-z](https://doi.org/10.1038/s41567-021-01287-z). URL: <https://doi.org/10.1038/s41567-021-01287-z>.
- [96] Jelena Mackeprang, Durga Bhaktavatsala Rao Dasari, and Jörg Wrachtrup. “A Reinforcement Learning approach for Quantum State Engineering”. In: *Quantum Machine Intelligence 2, 1(2020)* 2.1 (May 2019). DOI: [10.1007/s42484-020-00016-8](https://doi.org/10.1007/s42484-020-00016-8). arXiv: [1908.05981](https://arxiv.org/abs/1908.05981) [quant-ph].
- [97] Andrea Mari, Thomas R. Bromley, Josh Izaac, Maria Schuld, and Nathan Killoran. “Transfer learning in hybrid classical-quantum neural networks”. In: *Quantum* 4 (Oct. 2020), p. 340. ISSN: 2521-327X. DOI: [10.22331/q-2020-10-09-340](https://doi.org/10.22331/q-2020-10-09-340). URL: <http://dx.doi.org/10.22331/q-2020-10-09-340>.
- [98] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. “The theory of variational hybrid quantum-classical algorithms”. In: *New Journal of Physics* 18.2 (Feb. 2016), p. 023023. DOI: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023). URL: <https://dx.doi.org/10.1088/1367-2630/18/2/023023>.
- [99] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. “Barren plateaus in quantum neural network training landscapes”. In: *Nature Communications* 9.1 (Nov. 2018), p. 4812. ISSN: 2041-1723. DOI: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4). URL: <https://doi.org/10.1038/s41467-018-07090-4>.
- [100] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (June 1953), pp. 1087–1092. ISSN: 0021-9606. DOI: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114). URL: <https://doi.org/10.1063/1.1699114>.
- [101] Hrushikesh Mhaskar and Tomaso Poggio. *Deep vs. shallow networks : An approximation theory perspective*. 2016. arXiv: [1608.03287](https://arxiv.org/abs/1608.03287) [cs.LG]. URL: <https://arxiv.org/abs/1608.03287>.
- [102] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. “Asynchronous Methods for Deep Reinforcement Learning”. In: *Proceedings of The 33rd International Conference on Machine Learning*. Ed. by Maria Florina Balcan and Kilian Q. Weinberger. Vol. 48. Proceedings of Machine Learning Research. New York, New York, USA: PMLR, June 2016, pp. 1928–1937. URL: <https://proceedings.mlr.press/v48/mniha16.html>.

- [103] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. “Human-level control through deep reinforcement learning”. In: *Nature* 518.7540 (Feb. 2015), pp. 529–533. ISSN: 1476-4687. DOI: [10.1038/nature14236](https://doi.org/10.1038/nature14236). URL: <https://doi.org/10.1038/nature14236>.
- [104] Gordon Moore. “Cramming More Components onto Integrated Circuits (1965)”. In: *Ideas That Created the Future*. The MIT Press, Feb. 2021, pp. 261–266. ISBN: 9780262363174. DOI: [10.7551/mitpress/12274.003.0027](https://doi.org/10.7551/mitpress/12274.003.0027). URL: <http://dx.doi.org/10.7551/mitpress/12274.003.0027>.
- [105] Mikko Möttönen, Juha J. Vartiainen, Ville Bergholm, and Martti M. Salomaa. “Transformation of quantum states using uniformly controlled rotations”. In: *Quantum Info. Comput.* 5.6 (Sept. 2005), pp. 467–473. ISSN: 1533-7146. DOI: [10.5555/2011670.2011675](https://doi.org/10.5555/2011670.2011675).
- [106] Vinod Nair and Geoffrey E. Hinton. “Rectified linear units improve restricted boltzmann machines”. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML’10. Haifa, Israel: Omnipress, 2010, pp. 807–814. ISBN: 9781605589077. DOI: [10.5555/3104322.3104425](https://doi.org/10.5555/3104322.3104425).
- [107] Allen Newell and Herbert A. Simon. “Computer science as empirical inquiry: symbols and search”. In: *Commun. ACM* 19.3 (Mar. 1976), pp. 113–126. ISSN: 0001-0782. DOI: [10.1145/360018.360022](https://doi.org/10.1145/360018.360022). URL: <https://doi.org/10.1145/360018.360022>.
- [108] Alexander Quinn Nichol and Prafulla Dhariwal. “Improved Denoising Diffusion Probabilistic Models”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, July 2021, pp. 8162–8171. URL: <https://proceedings.mlr.press/v139/nichol21a.html>.
- [109] Michael A Nielsen. “A simple formula for the average gate fidelity of a quantum dynamical operation”. In: *Physics Letters A* 303.4 (2002), pp. 249–252. ISSN: 0375-9601. DOI: [10.1016/S0375-9601\(02\)01272-0](https://doi.org/10.1016/S0375-9601(02)01272-0). URL: <https://www.sciencedirect.com/science/article/pii/S0375960102012720>.
- [110] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, June 2012. ISBN: 9780511976667. DOI: [10.1017/cbo9780511976667](https://doi.org/10.1017/cbo9780511976667). URL: <http://dx.doi.org/10.1017/cbo9780511976667>.

-
- [111] Murphy Yuezhen Niu, Sergio Boixo, Vadim N. Smelyanskiy, and Hartmut Neven. “Universal quantum control through deep reinforcement learning”. In: *npj Quantum Information* 5.1 (Apr. 2019), p. 33. ISSN: 2056-6387. DOI: [10.1038/s41534-019-0141-3](https://doi.org/10.1038/s41534-019-0141-3). URL: <https://doi.org/10.1038/s41534-019-0141-3>.
- [112] Yash J. Patel, Akash Kundu, Mateusz Ostaszewski, Xavier Bonet-Monroig, Vedran Dunjko, and Onur Danaci. *Curriculum reinforcement learning for quantum architecture search under hardware errors*. 2024. arXiv: [2402 . 03500](https://arxiv.org/abs/2402.03500) [quant-ph]. URL: <https://arxiv.org/abs/2402.03500>.
- [113] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. “Data re-uploading for a universal quantum classifier”. In: *Quantum* 4 (Feb. 2020), p. 226. ISSN: 2521-327X. DOI: [10.22331/q-2020-02-06-226](https://doi.org/10.22331/q-2020-02-06-226). URL: <http://dx.doi.org/10.22331/q-2020-02-06-226>.
- [114] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. “A variational eigenvalue solver on a photonic quantum processor”. In: *Nature Communications* 5.1 (July 2014), p. 4213. ISSN: 2041-1723. DOI: [10 . 1038 / ncomms5213](https://doi.org/10.1038/ncomms5213). URL: <https://doi.org/10.1038/ncomms5213>.
- [115] Karl Raimund Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. London, England, 1962. DOI: [10.2307/2551371](https://doi.org/10.2307/2551371).
- [116] David Powers. “Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation”. English. In: *Journal of Machine Learning Technologies* 2.1 (2011), pp. 37–63. ISSN: 2229-3981.
- [117] Matteo G. Pozzi, Steven J. Herbert, Akash Sengupta, and Robert D. Mullins. “Using Reinforcement Learning to Perform Qubit Routing in Quantum Compilers”. In: *ACM Transactions on Quantum Computing* 3.2 (May 2022). DOI: [10.1145/3520434](https://doi.org/10.1145/3520434). URL: <https://doi.org/10.1145/3520434>.
- [118] John Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (Aug. 2018), p. 79. ISSN: 2521-327X. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79). URL: <https://doi.org/10.22331/q-2018-08-06-79>.
- [119] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. “High-Resolution Image Synthesis with Latent Diffusion Models”. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022, pp. 10674–10685. DOI: [10.1109/CVPR52688.2022.01042](https://doi.org/10.1109/CVPR52688.2022.01042).

- [120] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. “Learning representations by back-propagating errors”. In: *Nature* 323.6088 (Oct. 1986), pp. 533–536. ISSN: 1476-4687. DOI: [10.1038/323533a0](https://doi.org/10.1038/323533a0). URL: <http://dx.doi.org/10.1038/323533a0>.
- [121] Stuart Russell and Peter Norvig. *Artificial Intelligence, Global Edition A Modern Approach*. Pearson Deutschland, 2021, p. 1168. ISBN: 9781292401133. URL: <https://elibrary.pearson.de/book/99.150005/9781292401171>.
- [122] Antoine de Saint-Exupéry. *Wind, Sand and Stars*. Trans. by Lewis Galantière. Cambridge University Press (CUP), Mar. 1939, pp. 314–314. DOI: [10.1017/s0368393100105085](https://doi.org/10.1017/s0368393100105085).
- [123] Norbert Schuch and Jens Siewert. “Natural two-qubit gate for quantum computation using the XY interaction”. In: *Phys. Rev. A* 67 (3 Mar. 2003), p. 032301. DOI: [10.1103/PhysRevA.67.032301](https://doi.org/10.1103/PhysRevA.67.032301). URL: <https://link.aps.org/doi/10.1103/PhysRevA.67.032301>.
- [124] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. “Evaluating analytic gradients on quantum hardware”. In: *Phys. Rev. A* 99 (3 Mar. 2019), p. 032331. DOI: [10.1103/PhysRevA.99.032331](https://doi.org/10.1103/PhysRevA.99.032331). URL: <https://link.aps.org/doi/10.1103/PhysRevA.99.032331>.
- [125] Maria Schuld, Alex Bocharov, Krysta M. Svore, and Nathan Wiebe. “Circuit-centric quantum classifiers”. In: *Physical Review A* 101.3 (Mar. 2020). ISSN: 2469-9934. DOI: [10.1103/physreva.101.032308](https://doi.org/10.1103/physreva.101.032308). URL: <http://dx.doi.org/10.1103/PhysRevA.101.032308>.
- [126] Maria Schuld and Francesco Petruccione. *Supervised Learning with Quantum Computers*. Springer International Publishing, 2018. ISBN: 9783319964249. DOI: [10.1007/978-3-319-96424-9](https://doi.org/10.1007/978-3-319-96424-9). URL: <http://dx.doi.org/10.1007/978-3-319-96424-9>.
- [127] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. *High-Dimensional Continuous Control Using Generalized Advantage Estimation*. 2018. arXiv: [1506.02438](https://arxiv.org/abs/1506.02438) [cs.LG]. URL: <https://arxiv.org/abs/1506.02438>.
- [128] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. *Proximal Policy Optimization Algorithms*. 2017. arXiv: [1707.06347](https://arxiv.org/abs/1707.06347) [cs.LG]. URL: <https://arxiv.org/abs/1707.06347>.
- [129] Jaime Sevilla, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos. “Compute Trends Across Three Eras of Machine Learning”. In: *2022 International Joint Conference on Neural Networks (IJCNN)*. 2022, pp. 1–8. DOI: [10.1109/IJCNN55064.2022.9891914](https://doi.org/10.1109/IJCNN55064.2022.9891914).

- [130] V.V. Shende, S.S. Bullock, and I.L. Markov. “Synthesis of quantum-logic circuits”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 25.6 (2006), pp. 1000–1010. DOI: [10.1109/TCAD.2005.855930](https://doi.org/10.1109/TCAD.2005.855930).
- [131] P.W. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 124–134. DOI: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700).
- [132] P.W. Shor. “Fault-tolerant quantum computation”. In: *Proceedings of 37th Conference on Foundations of Computer Science*. 1996, pp. 56–65. DOI: [10.1109/SFCS.1996.548464](https://doi.org/10.1109/SFCS.1996.548464).
- [133] Peter W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493–R2496. DOI: [10.1103/PhysRevA.52.R2493](https://doi.org/10.1103/PhysRevA.52.R2493). URL: <https://link.aps.org/doi/10.1103/PhysRevA.52.R2493>.
- [134] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. “Mastering the game of Go with deep neural networks and tree search”. In: *Nature* 529.7587 (Jan. 2016), pp. 484–489. ISSN: 1476-4687. DOI: [10.1038/nature16961](https://doi.org/10.1038/nature16961). URL: <https://doi.org/10.1038/nature16961>.
- [135] Sukin Sim, Peter D. Johnson, and Alán Aspuru-Guzik. “Expressibility and Entangling Capability of Parameterized Quantum Circuits for Hybrid Quantum-Classical Algorithms”. In: *Advanced Quantum Technologies* 2.12 (2019), p. 1900070. DOI: [10.1002/qute.201900070](https://doi.org/10.1002/qute.201900070). eprint: <https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/qute.201900070>. URL: <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/qute.201900070>.
- [136] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. “t|ket>: A retargetable compiler for NISQ devices”. English. In: *Quantum Science and Technology* 6.1 (Nov. 2020). DOI: [10.1088/2058-9565/ab8e92](https://doi.org/10.1088/2058-9565/ab8e92).
- [137] Andrea Skolik, Sofiene Jerbi, and Vedran Dunjko. “Quantum agents in the Gym: a variational quantum algorithm for deep Q-learning”. en. In: *Quantum* 6 (May 2022), p. 720. DOI: [10.22331/q-2022-05-24-720](https://doi.org/10.22331/q-2022-05-24-720). URL: <http://dx.doi.org/10.22331/q-2022-05-24-720>.
- [138] Jiaming Song, Chenlin Meng, and Stefano Ermon. *Denoising Diffusion Implicit Models*. 2022. arXiv: [2010.02502](https://arxiv.org/abs/2010.02502) [cs.LG]. URL: <https://arxiv.org/abs/2010.02502>.

- [139] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. “Consistency models”. In: *Proceedings of the 40th International Conference on Machine Learning*. ICML’23. Honolulu, Hawaii, USA: JMLR.org, 2023. DOI: [10.5555/3618408.3619743](https://doi.org/10.5555/3618408.3619743).
- [140] Anders Sørensen and Klaus Mølmer. “Entanglement and quantum computation with ions in thermal motion”. In: *Phys. Rev. A* 62 (2 July 2000), p. 022311. DOI: [10.1103/PhysRevA.62.022311](https://doi.org/10.1103/PhysRevA.62.022311). URL: <https://link.aps.org/doi/10.1103/PhysRevA.62.022311>.
- [141] Gerhard Stenzel, Sebastian Zielinski, Michael Kölle, Philipp Altmann, Jonas Nüßlein, and Thomas Gabor. “Qandle: Accelerating State Vector Simulation Using Gate-Matrix Caching and Circuit Splitting”. In: *Proceedings of the 17th International Conference on Agents and Artificial Intelligence - Volume 1: QAIO*. INSTICC. SciTePress, 2025, pp. 715–723. ISBN: 978-989-758-737-5. DOI: [10.5220/0013343500003890](https://doi.org/10.5220/0013343500003890).
- [142] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. “On the importance of initialization and momentum in deep learning”. In: *Proceedings of the 30th International Conference on Machine Learning*. Ed. by Sanjoy Dasgupta and David McAllester. Vol. 28. Proceedings of Machine Learning Research 3. Atlanta, Georgia, USA: PMLR, June 2013, pp. 1139–1147. URL: <https://proceedings.mlr.press/v28/sutskever13.html>.
- [143] R.S. Sutton and A.G. Barto. “Reinforcement Learning: An Introduction”. In: *IEEE Transactions on Neural Networks* 9.5 (1998), pp. 1054–1054. DOI: [10.1109/TNN.1998.712192](https://doi.org/10.1109/TNN.1998.712192).
- [144] Richard S. Sutton. “Learning to predict by the methods of temporal differences”. In: *Machine Learning* 3.1 (Aug. 1988), pp. 9–44. ISSN: 1573-0565. DOI: [10.1007/BF00115009](https://doi.org/10.1007/BF00115009). URL: <https://doi.org/10.1007/BF00115009>.
- [145] Richard S. Sutton and Andrew G. Barto. “Toward a modern theory of adaptive networks: Expectation and prediction.” In: *Psychological Review* 88.2 (1981), pp. 135–170. DOI: [10.1037/0033-295X.88.2.135](https://doi.org/10.1037/0033-295X.88.2.135). URL: <https://doi.org/10.1037/0033-295X.88.2.135>.
- [146] Richard S. Sutton, David McAllester, Satinder Singh, and Yishay Mansour. “Policy gradient methods for reinforcement learning with function approximation”. In: *Proceedings of the 13th International Conference on Neural Information Processing Systems*. NIPS’99. Denver, CO: MIT Press, 1999, pp. 1057–1063.
- [147] Wei Tang, Yiheng Duan, Yaroslav Kharkov, Rasool Fakoor, Eric Kessler, and Yunong Shi. “AlphaRouter: Quantum Circuit Routing with Reinforcement Learning and Tree Search”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Vol. 01. 2024, pp. 930–940. DOI: [10.1109/QCE60285.2024.00112](https://doi.org/10.1109/QCE60285.2024.00112).

- [148] Barbara M. Terhal. “Quantum error correction for quantum memories”. en. In: *Reviews of Modern Physics* 87.2 (Apr. 2015), pp. 307–346. DOI: [10.1103/revmodphys.87.307](https://doi.org/10.1103/revmodphys.87.307). URL: <http://dx.doi.org/10.1103/RevModPhys.87.307>.
- [149] *The Ultimate Quotable Einstein*. Princeton University Press, 2011. ISBN: 9780691160146. URL: <http://www.jstor.org/stable/j.ctvs32rgr> (visited on 09/07/2025).
- [150] Tijmen Tieleman and Geoffrey Hinton. *Lecture 6.5—RMSPProp: Divide the gradient by a running average of its recent magnitude*. COURSERA: Neural Networks for Machine Learning. Accessed on 08/30/2025. 2012. URL: https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.
- [151] Alan Turing. “Computing Machinery and Intelligence”. In: *Mind* 59.236 (1950), pp. 433–60. DOI: [10.1093/mind/lix.236.433](https://doi.org/10.1093/mind/lix.236.433).
- [152] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. “Attention is all you need”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS’17. Long Beach, California, USA: Curran Associates Inc., 2017, pp. 6000–6010. ISBN: 9781510860964. DOI: [10.5555/3295222.3295349](https://doi.org/10.5555/3295222.3295349).
- [153] Pierre-François Verhulst. “Notice sur la loi que la population suit dans son accroissement”. In: *Correspondance mathématique et physique* 10 (1838), pp. 113–121.
- [154] Christopher J. C. H. Watkins and Peter Dayan. “Q-learning”. In: *Machine Learning* 8.3 (May 1992), pp. 279–292. ISSN: 1573-0565. DOI: [10.1007/BF00992698](https://doi.org/10.1007/BF00992698). URL: <https://doi.org/10.1007/BF00992698>.
- [155] William Zeng and Bob Coecke. “Quantum Algorithms for Compositional Natural Language Processing”. In: *Electronic Proceedings in Theoretical Computer Science* 221 (Aug. 2016), pp. 67–75. ISSN: 2075-2180. DOI: [10.4204/eptcs.221.8](https://doi.org/10.4204/eptcs.221.8). URL: <http://dx.doi.org/10.4204/EPTCS.221.8>.
- [156] Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. “Differentiable quantum architecture search”. In: *Quantum Science and Technology* 7.4 (Aug. 2022), p. 045023. DOI: [10.1088/2058-9565/ac87cd](https://doi.org/10.1088/2058-9565/ac87cd). URL: <http://dx.doi.org/10.1088/2058-9565/ac87cd>.
- [157] Jun Zhu. “Synthetic data generation by diffusion models”. In: *National Science Review* 11.8 (Aug. 2024), nwae276. ISSN: 2095-5138. DOI: [10.1093/nsr/nwae276](https://doi.org/10.1093/nsr/nwae276). eprint: <https://academic.oup.com/nsr/article-pdf/11/8/nwae276/59743306/nwae276.pdf>. URL: <https://doi.org/10.1093/nsr/nwae276>.

- [158] Weiwei Zhu, Jiangtao Pi, and Qiuyuan Peng. “A Brief Survey of Quantum Architecture Search”. In: *Proceedings of the 6th International Conference on Algorithms, Computing and Systems*. ICACS '22. Larissa, Greece: Association for Computing Machinery, 2023. ISBN: 9781450397407. DOI: [10 . 1145 / 3564982 . 3564989](https://doi.org/10.1145/3564982.3564989). URL: <https://doi.org/10.1145/3564982.3564989>.
- [159] Xianchao Zhu and Xiaokai Hou. “Quantum architecture search via truly proximal policy optimization”. In: *Scientific Reports* 13.1 (Mar. 2023), p. 5157. ISSN: 2045-2322. DOI: [10 . 1038 / s41598 - 023 - 32349 - 2](https://doi.org/10.1038/s41598-023-32349-2). URL: <https://doi.org/10.1038/s41598-023-32349-2>.

A. Publications

In the appendix, we present all 16 papers forming the core findings of this thesis. A concise overview of these papers is provided in appendix A.0. For ease of reference, table 1.1 is reprinted below in table A.1. Finally, the full texts of all papers, including their core contributions and the corresponding credit in accordance with the regulations of LMU Munich, are appended.

Table A.1.: An overview of the main technical contents of this thesis, research questions and publications. Reprinted from table 1.1.

Thesis chapter	Question	Publications
3. Training Optimization Techniques	Q1	75, 76, 78–80, 87
	Q2	75, 76, 78, 84
4. Automated Circuit Design and Synthesis	Q3	73, 81–83
5. Quantum Reinforcement Learning	Q4	77, 84, 86, 87
6. Advanced Models and Applications	Q5	72, 79, 80, 85
	Q6	74

Use of AI tools The use of AI tools, including large language models, was strictly limited to polishing and copy-editing purposes, i.e., refining language for clarity and style. All scientific content, research ideas, analyses, experimental designs, and conclusions presented in this thesis remain entirely the result of the author’s independent work and intellectual contribution. All substantive contributions conceived and developed in the publications subject to this thesis are summarized and appropriately credited in the following chapter.

A.0. Overview

The papers are listed with their number within the 16 papers, their reference number in the bibliography and the chapter where they are discussed. The full texts can be found on the specified pages.

Sec.	Ref.	Ch.	Publication	Page
A.1	[72]	6	Michael Kölle, Afrae Ahouzi, Pascal Debus, Robert Müller, Daniëlle Schuman, and Claudia Linnhoff-Popien. “Towards Efficient Quantum Anomaly Detection: One-Class SVMs Using Variable Subsampling and Randomized Measurements”. In: <i>Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART</i> . INSTICC. SciTePress, 2024, pp. 324–335	88
A.2	[73]	4	Michael Kölle, Tom Bintener, Maximilian Zorn, Gerhard Stenzel, Leo Sünkel, Thomas Gabor, and Claudia Linnhoff-Popien. “Evaluating Mutation Techniques in Genetic Algorithm-Based Quantum Circuit Synthesis”. In: <i>Proceedings of the Genetic and Evolutionary Computation Conference</i> . Association for Computing Machinery, July 2025, pp. 907–915	89
A.3	[74]	6	Michael Kölle, Alexander Feist, Jonas Stein, Sebastian Wölckert, and Claudia Linnhoff-Popien. “Evaluating Parameter-Based Training Performance of Neural Networks and Variational Quantum Circuits”. In: <i>Computational Science – ICCS 2025</i> . Ed. by Michael H. Lees, Wentong Cai, Siew Ann Cheong, Yi Su, David Abramson, Jack J. Dongarra, and Peter M. A. Sloot. Springer Nature Switzerland, 2025, pp. 275–282	90
A.4	[75]	3	Michael Kölle, Alessandro Giovagnoli, Jonas Stein, Maximilian Mansky, Julian Hager, and Claudia Linnhoff-Popien. “Improving Convergence for Quantum Variational Classifiers Using Weight Re-Mapping”. In: <i>Proceedings of the 15th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART</i> . INSTICC. SciTePress, 2023, pp. 251–258	91

Continued on next page

Sec.	Ref.	Ch.	Publication	Page
A.5	[76]	3	Michael Kölle, Alessandro Giovagnoli, Jonas Stein, Maximilian Balthasar Mansky, Julian Hager, Tobias Rohe, Robert Müller, and Claudia Linnhoff-Popien. “Weight Re-mapping for Variational Quantum Algorithms”. In: <i>Agents and Artificial Intelligence: 15th International Conference, ICAART 2023, Lisbon, Portugal, February 22–24, 2023, Revised Selected Papers</i> . Springer-Verlag, 2024, pp. 286–309	92
A.6	[77]	5	Michael Kölle, Mohamad Hgog, Fabian Ritz, Philipp Altmann, Maximilian Zorn, Jonas Stein, and Claudia Linnhoff-Popien. “Quantum Advantage Actor-Critic for Reinforcement Learning”. In: <i>Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART</i> . INSTICC. SciTePress, 2024, pp. 297–304	93
A.7	[78]	3	Michael Kölle, Leonhard Klingert, Julian Schönberger, Philipp Altmann, Tobias Rohe, and Claudia Linnhoff-Popien. “Investigating the Lottery Ticket Hypothesis for Variational Quantum Circuits”. In: <i>Proceedings of the 18th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART</i> . INSTICC. SciTePress, Mar. 2026, pp. 2433–2440	94
A.8	[79]	6	Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, Julian Hager, Sebastian Zielinski, and Claudia Linnhoff-Popien. “Quantum Transfer Learning: The Impact of Classical Preprocessing”. In: ed. by Ana Paula Rocha, Luc Steels, and Jaap van den Herik. Springer Nature Switzerland, Apr. 2025, pp. 368–403	95
A.9	[80]	6	Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, and Claudia Linnhoff-Popien. “Disentangling Quantum and Classical Contributions in Hybrid Quantum Machine Learning Architectures”. In: <i>Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART</i> . INSTICC. SciTePress, 2024, pp. 649–656	96

Continued on next page

Sec.	Ref.	Ch.	Publication	Page
A.10	[81]	4	Michael Kölle, Simon Salfer, Tobias Rohe, Philipp Altmann, and Claudia Linnhoff-Popien. “Quantum Architecture Search for Solving Quantum Machine Learning Tasks”. In: <i>Proceedings of the 18th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART</i> . INSTICC. SciTePress, Mar. 2026, pp. 1254–1263	97
A.11	[82]	4	Michael Kölle, Karola Schneider, Sabrina Egger, Felix Topp, Thomy Phan, Philipp Altmann, Jonas Nüßlein, and Claudia Linnhoff-Popien. “Architectural Influence on Variational Quantum Circuits in Multi-Agent Reinforcement Learning: Evolutionary Strategies for Optimization”. In: ed. by Ana Paula Rocha, Luc Steels, and Jaap van den Herik. Springer Nature Switzerland, Apr. 2025, pp. 50–79	98
A.12	[83]	4	Michael Kölle, Tom Schubert, Philipp Altmann, Maximilian Zorn, Jonas Stein, and Claudia Linnhoff-Popien. “A Reinforcement Learning Environment for Directed Quantum Circuit Synthesis”. In: <i>Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART</i> . INSTICC. SciTePress, 2024, pp. 83–94	99
A.13	[84]	5	Michael Kölle, Daniel Seidl, Maximilian Zorn, Philipp Altmann, Jonas Stein, and Thomas Gabor. “Optimizing Variational Quantum Circuits Using Metaheuristic Strategies in Reinforcement Learning”. In: <i>2024 IEEE International Conference on Quantum Computing and Engineering (QCE)</i> . IEEE Computer Society, Sept. 2024, pp. 323–328	100
A.14	[85]	6	Michael Kölle, Gerhard Stenzel, Jonas Stein, Sebastian Zielinski, Bjorn Ommer, and Claudia Linnhoff-Popien. “Quantum Denoising Diffusion Models”. In: <i>2024 IEEE International Conference on Quantum Software (QSW)</i> . IEEE Computer Society, July 2024, pp. 88–98	101

Continued on next page

Sec.	Ref.	Ch.	Publication	Page
A.15	[86]	5	Michael Kölle, Felix Topp, Thomy Phan, Philipp Altmann, Jonas Nüßlein, and Claudia Linnhoff-Popien. “Multi-Agent Quantum Reinforcement Learning Using Evolutionary Optimization”. In: <i>Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART</i> . INSTICC. SciTePress, 2024, pp. 71–82	102
A.16	[87]	3	Michael Kölle, Timo Witter, Tobias Rohe, Gerhard Stenzel, Philipp Altmann, and Thomas Gabor. “A Study on Optimization Techniques for Variational Quantum Circuits in Reinforcement Learning”. In: <i>2024 IEEE International Conference on Quantum Software (QSW)</i> . IEEE Computer Society, July 2024, pp. 157–167	103

A.1. Towards Efficient Quantum Anomaly Detection: One-Class SVMs Using Variable Subsampling and Randomized Measurements

Publication

Michael Kölle, Afrae Ahouzi, Pascal Debus, Robert Müller, Daniëlle Schuman, and Claudia Linnhoff-Popien. “Towards Efficient Quantum Anomaly Detection: One-Class SVMs Using Variable Subsampling and Randomized Measurements”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*. INSTICC. SciTePress, 2024, pp. 324–335

DOI: <https://doi.org/10.5220/0012381200003636>

URL: <https://www.scitepress.org/Papers/2024/123812/>

CODE: <https://github.com/AfraeA/q-anomaly>

Contribution

Focuses on scaling quantum anomaly detection using a one-class Support Vector Machine with quantum kernels. Proposes randomized measurements and variable subsampling to reduce time complexity from quadratic to near-linear, improving practicality while retaining strong average precision.

Credit

Kölle consulted on the core ideas and directed the overall research. Ahouzi performed experiments under Kölle’s guidance as part of her Bachelor’s thesis. Debus, Müller, Schuman discussed the approach, provided various ideas and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Core contribution to chapter 6 regarding the research question Q5 from section 1.2.

A.2. Evaluating Mutation Techniques in Genetic Algorithm-Based Quantum Circuit Synthesis

Publication

Michael Kölle, Tom Bintener, Maximilian Zorn, Gerhard Stenzel, Leo Sünkel, Thomas Gabor, and Claudia Linnhoff-Popien. “Evaluating Mutation Techniques in Genetic Algorithm-Based Quantum Circuit Synthesis”. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. Association for Computing Machinery, July 2025, pp. 907–915

DOI: <https://doi.org/10.1145/3712256.3726402>

URL: <https://doi.org/10.1145/3712256.3726402>

Contribution

Analyzes the influence of different mutation operators in genetic algorithm-based quantum circuit synthesis. Identifies a combination of addition, deletion and swap mutations as most effective in balancing circuit fidelity, depth and T operation usage.

Credit

Kölle conceived the idea, developed the research framework and provided comprehensive guidance throughout the project. Bintener conducted experiments under Kölle’s direction as part of his Bachelor’s thesis. Zorn, Stenzel, Sünkel, Gabor provided discussions, ideas and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Supporting contribution to chapter 4 regarding the research question Q3 from section 1.2.

A.3. Evaluating Parameter-Based Training Performance of Neural Networks and Variational Quantum Circuits

Publication

Michael Kölle, Alexander Feist, Jonas Stein, Sebastian Wölckert, and Claudia Linnhoff-Popien. “Evaluating Parameter-Based Training Performance of Neural Networks and Variational Quantum Circuits”. In: *Computational Science – ICCS 2025*. Ed. by Michael H. Lees, Wentong Cai, Siew Ann Cheong, Yi Su, David Abramson, Jack J. Dongarra, and Peter M. A. Sloot. Springer Nature Switzerland, 2025, pp. 275–282

DOI: https://doi.org/10.1007/978-3-031-97635-3_33

CODE: <https://github.com/alexander-feist/nn-vqc-params>

Contribution

Compares neural networks and variational quantum circuits across supervised and reinforcement learning tasks, focusing on varying parameter counts and associated training times. Shows that while variational quantum circuits can achieve performance on par with neural networks using significantly fewer parameters, they currently require substantially longer training durations—underscoring their potential efficiency advantages as quantum hardware matures.

Credit

Kölle initiated the project, developed the conceptual framework and supervised the research. Feist carried out the experiments as part of his Bachelor’s thesis under Kölle’s guidance. Wölckert and Stein discussed the methods and reviewed the results. Linnhoff-Popien provided strategic oversight and reviewed the results.

Purpose

Main focus of chapter 6 regarding the research question Q6 from section 1.2.

A.4. Improving Convergence for Quantum Variational Classifiers Using Weight Re-Mapping

Publication

Michael Kölle, Alessandro Giovagnoli, Jonas Stein, Maximilian Mansky, Julian Hager, and Claudia Linnhoff-Popien. “Improving Convergence for Quantum Variational Classifiers Using Weight Re-Mapping”. In: *Proceedings of the 15th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*. INSTICC. SciTePress, 2023, pp. 251–258

DOI: <https://doi.org/10.5220/0011696300003393>

URL: <https://www.scitepress.org/Papers/2023/116963/>

CODE: <https://github.com/michaelkoelle/qw-map>

Contribution

Introduces weight re-mapping for variational quantum circuits to resolve the mismatch between trainable parameters and rotational angles. Empirically shows that weight re-mapping can improve convergence and test accuracy in quantum classifiers for specific tasks.

Credit

Kölle originated the weight re-mapping concept, developed the theoretical foundation addressing variational quantum circuit limitations and provided research direction. Giovagnoli implemented the methodology and conducted experiments under Kölle’s guidance. Stein, Mansky, Hager discussed the approach, contributed to the background section and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Core contribution to chapter 3 regarding research questions Q1 and Q2 from section 1.2.

A.5. Weight Re-mapping for Variational Quantum Algorithms

Publication

Michael Kölle, Alessandro Giovagnoli, Jonas Stein, Maximilian Balthasar Mansky, Julian Hager, Tobias Rohe, Robert Müller, and Claudia Linnhoff-Popien. “Weight Re-mapping for Variational Quantum Algorithms”. In: *Agents and Artificial Intelligence: 15th International Conference, ICAART 2023, Lisbon, Portugal, February 22–24, 2023, Revised Selected Papers*. Springer-Verlag, 2024, pp. 286–309

DOI: https://doi.org/10.1007/978-3-031-55326-4_14

URL: https://doi.org/10.1007/978-3-031-55326-4_14

CODE: <https://github.com/michaelkoelle/qw-map>

Contribution

This is a substantially extended version of Kölle et al. [75], presenting a broader set of weight re-mapping functions, more comprehensive experimental datasets and a deeper analysis of their impact on convergence speed and overall classification accuracy in variational quantum circuits.

Credit

Kölle led this extension of previous work, conceptualized the expanded framework, designed the experimental protocol and conducted all extended experimentation and analysis. Giovagnoli, Stein, Mansky, Hager, Müller and Rohe contributed additional discussions, conceptual refinements and helped review the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Core contribution to chapter 3 regarding research questions Q1 and Q2 from section 1.2.

A.6. Quantum Advantage Actor-Critic for Reinforcement Learning

Publication

Michael Kölle, Mohamad Hgog, Fabian Ritz, Philipp Altmann, Maximilian Zorn, Jonas Stein, and Claudia Linnhoff-Popien. “Quantum Advantage Actor-Critic for Reinforcement Learning”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART*. INSTICC. SciTePress, 2024, pp. 297–304

Michael Kölle and Mohamad Hgog contributed equally to this work

DOI: <https://doi.org/10.5220/0012383900003636>

URL: <https://www.scitepress.org/Papers/2024/123839/>

CODE: <https://github.com/Hajuj/Quantum-A2C>

Contribution

Introduces a hybrid quantum-classical actor-critic framework by replacing select parts of the Advantage Actor-Critic pipeline with variational quantum circuits. Demonstrates performance gains in Cart Pole, particularly when only the actor or the critic is quantum.

Credit

Kölle formulated the main idea, initial concepts and provided research guidance. Hgog conducted experiments under Kölle’s supervision as part of his Bachelor’s thesis. Altmann, Ritz, Zorn, Stein offered additional discussions, ideas and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Core contribution to chapter 5 regarding the research question Q4 from section 1.2.

A.7. Investigating the Lottery Ticket Hypothesis for Variational Quantum Circuits

Publication

Michael Kölle, Leonhard Klingert, Julian Schönberger, Philipp Altmann, Tobias Rohe, and Claudia Linnhoff-Popien. “Investigating the Lottery Ticket Hypothesis for Variational Quantum Circuits”. In: *Proceedings of the 18th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART*. INSTICC. SciTePress, Mar. 2026, pp. 2433–2440

DOI: <https://doi.org/10.5220/0014247600004052>

URL: <https://www.scitepress.org/Link.aspx?doi=10.5220/0014247600004052>

CODE: <https://github.com/916750/lottery-ticket-vqc>

Contribution

Investigates the applicability of the Lottery Ticket Hypothesis to Variational Quantum Circuits. Empirically demonstrates that the weak LTH holds for VQCs, identifying “winning tickets” with as few as 26% of the original parameters. Shows that pruning can serve as an effective strategy to mitigate the barren plateau phenomenon.

Credit

Kölle conceived the core concept and supervised the project. Klingert conducted the experiments as part of his Bachelor’s thesis under Kölle’s guidance. Schönberger, Altmann, Mansky and Linnhoff-Popien reviewed the results and provided feedback.

Purpose

Main focus of chapter 3 regarding the research question Q2 from section 1.2.

A.8. Quantum Transfer Learning: The Impact of Classical Preprocessing

Publication

Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, Julian Hager, Sebastian Zielinski, and Claudia Linnhoff-Popien. “Quantum Transfer Learning: The Impact of Classical Preprocessing”. In: ed. by Ana Paula Rocha, Luc Steels, and Jaap van den Herik. Springer Nature Switzerland, Apr. 2025, pp. 368–403

DOI: https://doi.org/10.1007/978-3-031-87327-0_18

CODE: <https://github.com/javajonny/AE-and-VQC>

Contribution

This is a substantially extended version of Kölle et al. [80]. It introduces a methodologically extended hybrid architecture using a classical autoencoder for unsupervised dimensionality reduction to fairly evaluate the performance contributions of the classical and quantum components.

Credit

Kölle developed the core idea and provided research supervision. Maurer conducted the experiments as part of his Bachelor’s thesis under Kölle’s guidance. Altmann, Sünkel, Stein, Hager, Zielinski and Linnhoff-Popien provided discussions and reviewed the results.

Purpose

Main focus of chapter 6 regarding the research questions Q1 and Q5 from section 1.2.

A.9. Disentangling Quantum and Classical Contributions in Hybrid Quantum Machine Learning Architectures

Publication

Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, and Claudia Linnhoff-Popien. “Disentangling Quantum and Classical Contributions in Hybrid Quantum Machine Learning Architectures”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART*. INSTICC. SciTePress, 2024, pp. 649–656

DOI: <https://doi.org/10.5220/0012381600003636>

URL: <https://www.scitepress.org/Papers/2024/123816/>

CODE: <https://github.com/javajonny/AE-and-VQC>

Contribution

Presents a novel hybrid transfer learning architecture using an autoencoder for input compression followed by a variational quantum circuit, clarifying the classical and quantum contributions to the final classification performance.

Credit

Kölle developed the main idea and approach, established the theoretical framework for the analysis and provided research guidance throughout the project. Maurer implemented the methodology and performed experiments under Kölle’s supervision as part of his Bachelor’s thesis. Altmann, Sünkel, Stein provided conceptual discussions and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Main focus of chapter 6 regarding the research questions Q1 and Q5 from section 1.2.

A.10. Quantum Architecture Search for Solving Quantum Machine Learning Tasks

Publication

Michael Kölle, Simon Salfer, Tobias Rohe, Philipp Altmann, and Claudia Linnhoff-Popien. “Quantum Architecture Search for Solving Quantum Machine Learning Tasks”. In: *Proceedings of the 18th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*. INSTICC. SciTePress, Mar. 2026, pp. 1254–1263

DOI: <https://doi.org/10.5220/0014247700004052>

URL: <https://www.scitepress.org/Link.aspx?doi=10.5220/0014247700004052>

CODE: <https://github.com/916750/qas4ml>

Contribution

Introduces RL-QAS, a reinforcement learning framework to automate the discovery of variational quantum circuit architectures for machine learning tasks. Demonstrates on Iris and binary MNIST datasets that the agent autonomously finds low-complexity circuits with high test accuracy, outperforming standard manual designs.

Credit

Kölle conceived the idea, created the methodology and supervised the project. Salfer conducted the experiments as part of his Bachelor’s thesis under Kölle’s guidance. Rohe, Altmann and Linnhoff-Popien reviewed the results and provided feedback.

Purpose

Core contribution to chapter 4 regarding the research question Q3 from section 1.2.

A.11. Architectural Influence on Variational Quantum Circuits in Multi-Agent Reinforcement Learning: Evolutionary Strategies for Optimization

Publication

Michael Kölle, Karola Schneider, Sabrina Egger, Felix Topp, Thomy Phan, Philipp Altmann, Jonas Nüßlein, and Claudia Linnhoff-Popien. “Architectural Influence on Variational Quantum Circuits in Multi-Agent Reinforcement Learning: Evolutionary Strategies for Optimization”. In: ed. by Ana Paula Rocha, Luc Steels, and Jaap van den Herik. Springer Nature Switzerland, Apr. 2025, pp. 50–79

DOI: https://doi.org/10.1007/978-3-031-87327-0_3

CODE: <https://github.com/karolaschneider/vqc-opt-arch-evo>

Contribution

Extends the previous work Kölle et al. [86] on evolutionary multi-agent quantum reinforcement learning by focusing on architecture optimization. Proposes and compares three distinct evolutionary strategies (Gate-Based, Layer-Based and Prototype-Based) for co-evolving a VQC’s structure and parameters, finding that a flexible, mutation-only, gate-based approach performs best.

Credit

Kölle had the original idea, proposed the three evaluated strategies and supervised the research. Schneider conducted the experiments as part of her Bachelor’s thesis under Kölle’s supervision. Egger and Topp contributed to the preliminaries and related work sections. Phan, Altmann, Nüßlein and Linnhoff-Popien provided discussions and reviewed the results.

Purpose

Core contribution to chapter 4 regarding the research question Q3 from section 1.2.

A.12. A Reinforcement Learning Environment for Directed Quantum Circuit Synthesis

Publication

Michael Kölle, Tom Schubert, Philipp Altmann, Maximilian Zorn, Jonas Stein, and Claudia Linnhoff-Popien. “A Reinforcement Learning Environment for Directed Quantum Circuit Synthesis”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART*. INSTICC. SciTePress, 2024, pp. 83–94

DOI: <https://doi.org/10.5220/0012383200003636>

URL: <https://www.scitepress.org/Papers/2024/123832/>

CODE: <https://github.com/michaelkoelle/rl-qc-syn>

Contribution

Proposes a reinforcement learning environment to synthesize quantum circuits with the Clifford+T gate set, targeting specific quantum states. Evaluates solutions generated by Proximal Policy Optimization for 2-qubit Bell states and extends the environment to higher-qubit benchmarks.

Credit

Kölle conceived the main idea, established the theoretical framework, designed the specialized environment and provided research direction. Schubert conducted experiments under Kölle’s supervision as part of his Bachelor’s thesis. Altmann, Zorn, Stein contributed conceptual feedback and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Supporting contribution to chapter 4 regarding the research question Q3 from section 1.2.

A.13. Optimizing Variational Quantum Circuits Using Metaheuristic Strategies in Reinforcement Learning

Publication

Michael Kölle, Daniel Seidl, Maximilian Zorn, Philipp Altmann, Jonas Stein, and Thomas Gabor. “Optimizing Variational Quantum Circuits Using Metaheuristic Strategies in Reinforcement Learning”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE Computer Society, Sept. 2024, pp. 323–328

DOI: <https://doi.org/10.1109/QCE60285.2024.10300>

URL: <https://www.computer.org/csdl/proceedings-article/qce/2024/413702a323/23oqmUz6Jj0>

CODE: <https://github.com/michaelkoelle/metaheuristic-vis>

Contribution

Examines metaheuristic algorithms—Particle Swarm Optimization, Ant Colony Optimization, Tabu Search, Genetic Algorithm, Simulated Annealing and Harmony Search—for gradient-free quantum reinforcement learning. Demonstrates robust performance in MiniGrid tasks and Cart Pole, with notable success for Simulated Annealing and Particle Swarm Optimization.

Credit

Kölle established the research direction, designed the experimental methodology, conducted preliminary investigations and provided supervision throughout. Seidl implemented the methodologies and performed experiments under Kölle’s guidance as part of his Master’s thesis. Zorn, Altmann, Stein discussed the methods and reviewed the results. Gabor and Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Core contribution to chapter 3 and chapter 5 regarding research questions Q2 and Q4 from section 1.2.

A.14. Quantum Denoising Diffusion Models

Publication

Michael Kölle, Gerhard Stenzel, Jonas Stein, Sebastian Zielinski, Bjorn Ommer, and Claudia Linnhoff-Popien. “Quantum Denoising Diffusion Models”. In: *2024 IEEE International Conference on Quantum Software (QSW)*. IEEE Computer Society, July 2024, pp. 88–98

DOI: <https://doi.org/10.1109/QSW62656.2024.00023>

URL: <https://www.computer.org/csdl/proceedings-article/qsw/2024/684700a088/1ZMeeEZjPPy>

CODE: <https://github.com/michaelkoelle/quantum-diffusion>

Contribution

Proposes quantum denoising diffusion models to address slow sampling and large parameter counts in classical diffusion frameworks. Introduces two quantum diffusion architectures and a consistency-model-based single-step sampling design, surpassing classical baselines on MNIST, Fashion MNIST and CIFAR-10.

Credit

Kölle developed the quantum diffusion concept, executed initial experiments and directed the research. Stenzel conducted extended evaluations under Kölle’s guidance as part of his Master’s thesis. Stein, Zielinski discussed the methods, contributed ideas and reviewed the results. Ommer and Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Main focus of chapter 6 regarding the research question Q5 from section 1.2.

A.15. Multi-Agent Quantum Reinforcement Learning Using Evolutionary Optimization

Publication

Michael Kölle, Felix Topp, Thomy Phan, Philipp Altmann, Jonas Nüßlein, and Claudia Linnhoff-Popien. “Multi-Agent Quantum Reinforcement Learning Using Evolutionary Optimization”. In: *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART*. INSTICC. SciTePress, 2024, pp. 71–82

DOI: <https://doi.org/10.5220/0012382800003636>

URL: <https://www.scitepress.org/Papers/2024/123828>

CODE: <https://github.com/michaelkoelle/qmarl-evo>

Contribution

Extends quantum reinforcement learning to multi-agent settings by combining variational quantum circuits with evolutionary optimization. Demonstrates parameter efficiency in the Coin Game, approaching large classical network performance using only a fraction of its parameters.

Credit

Kölle conceived the main idea, devised the methodology and developed the experimental design. Topp implemented and executed experiments under Kölle’s guidance as part of his Bachelor’s thesis. Phan, Altmann, Nüßlein discussed the approach and reviewed the results. Linnhoff-Popien consulted the process and reviewed the results.

Purpose

Core contribution to chapter 5 regarding the research question Q4 from section 1.2.

A.16. A Study on Optimization Techniques for Variational Quantum Circuits in Reinforcement Learning

Publication

Michael Kölle, Timo Witter, Tobias Rohe, Gerhard Stenzel, Philipp Altmann, and Thomas Gabor. “A Study on Optimization Techniques for Variational Quantum Circuits in Reinforcement Learning”. In: *2024 IEEE International Conference on Quantum Software (QSW)*. IEEE Computer Society, July 2024, pp. 157–167

DOI: <https://doi.org/10.1109/QSW62656.2024.00031>

URL: <https://www.computer.org/csdl/proceedings-article/qsw/2024/684700a157/1ZMefLwAX1m>

CODE: <https://github.com/michaelkoelle/vqc-opt-rl>

Contribution

Investigates the impact of data re-uploading, input/output scaling and introduces exponential learning rate decay in variational quantum proximal policy optimization. Evaluates how these techniques stabilize hyperparameters, reduce parameter usage and bolster performance in Frozen Lake and Cart Pole tasks.

Credit

Kölle formulated the research approach, established the theoretical foundations, developed the investigation framework and provided direction throughout the project. Witter implemented methodologies and conducted experiments under Kölle’s supervision as part of his Bachelor’s thesis. Rohe, Stenzel, Altmann provided discussions, insights and reviewed the outcomes. Gabor consulted the process and reviewed the results.

Purpose

Core contribution to chapter 3 and chapter 5 regarding research questions Q1 and Q4 from section 1.2.