

# Quantum neural networks

Von der Fakultät für Mathematik und Physik  
der Gottfried Wilhelm Leibniz Universität Hannover

zur Erlangung des akademischen Grades

Doctor rerum naturalium

Dr. rer. nat.

genehmigte Dissertation von

**M. Sc. Kerstin Beer**

2022

**Examination board:**

Prof. Dr. Michèle Heurs (chair)

Prof. Dr. Tobias J. Osborne (supervisor)

Prof. Dr. Avishek Anand

**Referees:**

Prof. Dr. Tobias J. Osborne

Prof. Dr. Avishek Anand

Prof. Dr. Chiara Macchiavello

Day of the defence: 22.02.2022

*“I have never tried that before,  
so I think I should definitely be able to do that.”*  
Pippi Longstocking (Astrid Lindgren)

## Abstract

Quantum computing is one of the most exciting research areas of the last decades. At the same time, methods of machine learning have started to dominate science, industry and our everyday life. In this thesis we combine these two essential research topics of the 21st century and introduce *dissipative quantum neural networks* (DQNNs), which are designed for fully quantum learning tasks, are capable of universal quantum computation and have low memory requirements while training.

We start the discussion of this interdisciplinary topic by introducing artificial neural networks, which are a very common tool in classical machine learning. Next, we give an overview on quantum information. Here we focus on quantum algorithms and circuits, which are used to implement quantum neural networks. Moreover, we explain the opportunities and challenges arising with today's quantum computers.

The discussion of the architecture and training algorithm of the DQNNs forms the core of this work. These networks are optimised with training data pairs in form of input and desired output states and therefore can be used for characterising unknown or untrusted quantum devices. We not only demonstrate the generalisation behaviour of these quantum neural networks using classical simulations, but also implement them successfully on actual quantum computers.

To understand the ultimate limits for such quantum machine learning methods, we discuss the *quantum no free lunch* theorem, which describes a bound on the probability that a quantum device, which can be modelled as a unitary process and is optimised with quantum examples, gives an incorrect output for a random input. This gives us a tool to review the learning behaviour of quantum neural networks in general and the DQNNs in particular.

Moreover we expand the area of applications of DQNNs in two directions. In the first case, we include additional information beyond just the training data pairs: since quantum devices are always structured, the resulting data is always structured as well. We modify the DQNN's training algorithm such that knowledge about the graph-structure of the training data pairs is included in the training process and show that this can lead to better generalisation behaviour.

Both the original DQNN and the DQNN including graph structure are trained with data pairs in order to characterise an underlying relation. However, in the second extension of the algorithm we aim to learn characteristics of a set of quantum states in order to extend it to quantum states which have similar properties. Therefore we build a generative adversarial model where two DQNNs, called the generator and discriminator, are trained in a competitive way.

Overall, we observe that DQNNs can not only be trained efficiently but also, similar to their classical counterparts, modified to suit different applications.

*Keywords:* quantum computing, neural network, machine learning

## Kurzzusammenfassung

Quantencomputer bilden eines der spannendsten Forschungsgebiete der letzten Jahrzehnte. Zur gleichen Zeit haben Methoden des maschinellen Lernens begonnen die Wissenschaft, Industrie und unseren Alltag zu dominieren. In dieser Arbeit kombinieren wir diese beiden wichtigen Forschungsthemen des 21. Jahrhunderts und stellen *dissipative quantenneuronale Netze* (DQNNs) vor, die für Quantenlernaufgaben konzipiert sind, universelle Quantenberechnungen durchführen können und wenig Speicherbedarf beim Training benötigen.

Wir beginnen die Diskussion dieses interdisziplinären Themas mit der Einführung künstlicher neuronaler Netze, die beim klassischen maschinellen Lernen weit verbreitet sind. Dann geben wir einen Überblick über die Quanteninformationstheorie. Hier fokussieren wir uns auf die zur Implementierung von quantenneuronalen Netzen nötigen Quantenalgorithmen und -schaltungen. Außerdem erläutern wir die Chancen und Herausforderungen der heutigen Quantencomputer.

Die Diskussion der Architektur und des Trainingsalgorithmus der DQNNs bildet den Mittelpunkt dieser Arbeit. Diese Netzwerke werden mit Trainingsdatenpaaren in Form von Eingangs- und gewünschten Ausgangszuständen optimiert und können daher zur Charakterisierung unbekannter oder nicht vertrauenswürdiger Quantenbauelemente verwendet werden. Wir demonstrieren nicht nur das Generalisierungsverhalten dieser Netze anhand klassischer Simulationen, sondern konstruieren auch eine erfolgreiche Implementierung für Quantencomputer.

Um die ultimativen Grenzen solcher Methoden zum maschinellen Lernen von Quantendaten zu verstehen, führen wir das *quantum no free lunch*-Theorem ein, welches eine Begrenzung für die Wahrscheinlichkeit beschreibt, dass ein als unitärer Prozess modellierbares und mit Quantendaten optimiertes Quantenbauelement eine falsche Ausgabe für eine zufällige Eingabe herausgibt. Das Theorem gibt uns ein Werkzeug, um das Lernverhalten von quantenneuronalen Netzwerken im Allgemeinen und der DQNNs im Besonderen zu überprüfen.

Darüber hinaus erweitern wir den Anwendungsbereich von DQNNs auf zwei Weisen. Im ersten Fall beziehen wir Informationen zusätzlich zu den Trainingsdaten mit ein: Da Quantenbauelemente immer eine gewisse Struktur haben, sind auch die resultierenden Daten strukturiert. Wir modifizieren den Trainingsalgorithmus der DQNNs so, dass Kenntnisse über die Struktur genutzt werden können und zeigen, dass dies zu einem besseren Trainingsergebnis führen kann.

Sowohl das ursprüngliche DQNN als auch das Graphen-DQNN wird mit Datenpaaren trainiert, um eine zugrunde liegende Relation zu charakterisieren. Als zweite Erweiterung wollen wir jedoch die Eigenschaften einer Menge einzelner Quantenzustände untersuchen, um sie mit Quantenzuständen ähnlicher Eigenschaften zu erweitern. Daher konstruieren wir ein Modell, bei dem zwei DQNNs, Generator und Diskriminator genannt, kompetitiv trainiert werden.

Zusammenfassend stellen wir fest, dass DQNNs nicht nur effizient trainiert, sondern auch, ähnlich wie ihre klassischen Gegenstücke, an unterschiedliche Anwendungen angepasst werden können.

*Schlagwörter:* Quantencomputer, Neuronales Netz, Maschinelles Lernen

## Acknowledgements

Without question, I am very thankful that I was given the chance to be a Ph.D. student and write this thesis. This would have been inconceivable without the people named in the following.

Foremost, I would like to thank my supervisor Tobias J. Osborne for introducing me to the exciting field of quantum information, guiding me through the last years, and sharing his knowledge and experiences with me. I thank him for creating a well-organised, open-minded and safe working atmosphere, all his helpful advice, and making me sometimes feel a little bit like Pippi Longstocking, see page III.

Further, I would like to thank Avishek Anand for introducing me to the topic of deep neural networks in his lecture, Megha Khosla for various helpful discussions, and all people who cooperated with me on several exciting projects.

Many thanks go to all past and current members of the *quantum information group* in Hanover. I explicitly thank Kais Abdelkhalek, Thomas Cope, Friederike Dziemba, Terry Farrelly, Tobias Geib, Alexander Hahn, Wiebke Möller, Laura Niermann, Viktoria Schmiesing, Reinhard Werner, and Ramona Wolf for making our group a warm-hearted and amusing working place. Moreover, I thank Florian Oppermann for being the funniest and most generous IT administrator colleague one could wish for. I also thank all Master and Bachelor students I supervised for knowledge exchange and various enjoyable meetings in the last years.

I am very thankful for all the opportunities to have thrilling travel experiences and meet inspiring people. Particularly I would like to thank Felipe Montealegre-Mora for the most cheerful conferences. Moreover, I thank Hanover's university groups *Students for Future* and *Scientists for Future* for giving me the chance to learn and discuss science beyond quantum information.

Sometimes life can be hard, hence it is important to have good friends. I want to thank Bianka and Robert Heymann, Luisa Oyen, and all members of *Zirkus Johnass*, my second family. Further, I thank my friends at the acrobatic course of our university, and at the *Grüntöne Ensemble* for all the fun time we spend together. I thank my cello teacher Jacob Jordan for one certainly joyful hour every week. I also would like to thank Franziska Weeren for various long and understanding phone calls, as well as my therapist and her dog Erwin.

I would like to thank Thomas Cope, Tobias Geib, Dominik Lack, Gabriel Müller, Laura Niermann, Viktoria Schmiesing, Marvin Schwiering, Alexander Stottmeister, and Christian Struckmann for finding several typos in my drafts and, hence, helped me a lot to bring this thesis in its final form.

My biggest thanks go to my beloved family, my mum Renata Beer, my dad Christian Beer and my sister Anika Beer for their limitless support and love.

Most importantly, I want to thank my amazing partner, lovely husband and best friend Dominik Lack for being by my side for the last eleven years. I would like to thank him for his kind words and hugs whenever I panicked, for his endless love, and for all the snacks he brought me during the home office times. The last years, and this thesis would not have been possible without him.

---

## Publications

---

### Publications whose material appears in this thesis:

- K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf. Training deep quantum neural networks. *Nature Communications*, 11(1):1–6, 2020. doi:[10.1038/s41467-020-14454-2](https://doi.org/10.1038/s41467-020-14454-2).
- K. Poland, K. Beer, and T. J. Osborne. No free lunch for quantum machine learning, 2020, [arXiv:2003.14103](https://arxiv.org/abs/2003.14103).
- K. Beer, M. Khosla, J. Köhler, and T. J. Osborne. Quantum machine learning of graph-structured data. 2021, [arXiv:2103.10837](https://arxiv.org/abs/2103.10837).
- K. Beer, D. List, G. Müller, T. J. Osborne, and C. Struckmann. Training Quantum Neural Networks on NISQ Devices. 2021, [arXiv:2104.06081](https://arxiv.org/abs/2104.06081).
- K. Beer, and G. Müller. Dissipative quantum generative adversarial networks. 2021, [arXiv:2112.06088](https://arxiv.org/abs/2112.06088).

### Further publications:

- K. Beer and F. A. Dziemba. Phase-context decomposition of diagonal unitaries for higher-dimensional systems. *Physical Review A*, 2016. doi:[10.1103/PhysRevA.93.052333](https://doi.org/10.1103/PhysRevA.93.052333).
- K. Beer and T. J. Osborne. Contextuality and bundle diagrams. *Physical Review A*, 2018. doi:[10.1103/PhysRevA.98.052124](https://doi.org/10.1103/PhysRevA.98.052124).
- K. Beer, D. Bondarenko, A. Hahn, M. Kalabakov, N. Knust, L. Niermann, T. J. Osborne, C. Schridde, S. Seckmeyer, D. E. Stiegemann, and R. Wolf. From categories to anyons: a travelogue. 2018, [arXiv:1811.06670](https://arxiv.org/abs/1811.06670).

# Contents

Abstract . . . . .	IV
Kurzzusammenfassung . . . . .	V
Acknowledgements . . . . .	VII
Publications . . . . .	VIII
<b>1 Introduction</b>	<b>1</b>
<b>2 Classical neural networks</b>	<b>5</b>
2.1 Artificial neurons . . . . .	6
2.2 Network architecture . . . . .	9
2.3 Training data and loss functions . . . . .	13
2.4 Optimisation . . . . .	16
<b>3 Quantum information</b>	<b>21</b>
3.1 Quantum bits . . . . .	22
3.2 Quantum circuits . . . . .	28
3.3 Quantum algorithms . . . . .	32
3.4 Quantum computers . . . . .	35
3.5 Quantum neural networks . . . . .	36
<b>4 Dissipative quantum neural networks</b>	<b>39</b>
4.1 Network architecture . . . . .	40
4.2 Loss functions . . . . .	43
4.3 Training algorithm . . . . .	45
4.4 Universality . . . . .	56
4.5 Classical simulation . . . . .	58
4.6 NISQ device implementation . . . . .	61
4.7 Comparison to quantum approximate optimisation algorithm . . . . .	68



<b>5</b>	<b>No free lunch theorem</b>	<b>75</b>
5.1	Classical no free lunch theorem . . . . .	76
5.2	Quantum no free lunch theorem . . . . .	77
<b>6</b>	<b>Training with graph-structured quantum data</b>	<b>87</b>
6.1	Basic definitions . . . . .	88
6.2	Classical machine learning with graph-structured data . . . . .	89
6.3	Quantum graph-structured data . . . . .	91
6.4	Loss functions . . . . .	93
6.5	Training algorithm . . . . .	95
6.6	Classical simulation . . . . .	97
6.7	NISQ device implementation . . . . .	102
<b>7</b>	<b>Quantum generative adversarial networks</b>	<b>105</b>
7.1	Classical adversarial networks . . . . .	107
7.2	Dissipative quantum generative adversarial network . . . . .	111
7.3	Loss functions . . . . .	112
7.4	Training algorithm . . . . .	113
7.5	Classical simulation . . . . .	120
7.6	NISQ device implementation . . . . .	124
<b>8</b>	<b>Conclusion and outlook</b>	<b>127</b>
	<b>Appendix A Dissipative quantum neural networks</b>	<b>135</b>
	<b>Appendix B Training with graph-structured quantum data</b>	<b>143</b>
	<b>Appendix C Quantum generative adversarial networks</b>	<b>145</b>
	Bibliography . . . . .	150
	Curriculum vitae . . . . .	179



# 1

---

## Introduction

---

Results of *machine learning* (ML) [1–4], the well-known subfield of *artificial intelligence* where knowledge is gained from experiences rather than from instructions, have carried over into our everyday life in the last decade: our web search engines rearrange and optimise results based on learned user characteristics [5–7], we are used to traffic forecasting and always up-to-date commute-estimating apps [8, 9], spam and phishing emails can be automatically classified [10–13] and several apps even offer smart replies [14–16], on social media platforms the tools of ML allow recommendation of friends, posts or tags [17, 18], identification of illegal, unwanted or fake data [19–22] and actually allow the platforms to extract informations on the user’s personality, interests or mental health [23–27].

Also, generally, the application of ML techniques in neuroscience, medical diagnosis, and healthcare is widespread [28–31]. Here especially ML algorithms, which learn from data samples, are common. Such methods can be, for instance, used for side effect prediction of drugs [32] or for analysing or classifying tomography-computed, X-ray, and magnetic resonance images [33–40].

Since ML can be applied in nearly every area where enough data is accumulated, it is without question that also a multitude of industries make use of these mechanisms [41–43]. However, ML approaches can even be a powerful tool in reducing greenhouse gas emissions, helping society adapt to a changing climate and model extreme weather events [44–46].

Because ML is, as illustrated above, used on so many different problems in various areas, plenty of different methods exist. However, the use of *artificial neural networks* (NNs) [3, 4, 47] is very common. These networks consist of a layered architecture built up of fundamental units, originally inspired by the neurons of a human’s neural network. These units are linked via weighted connections, which

## 1. Introduction

are updated during a training process. After that, the network ideally is able to master specific tasks, such as classifying pictures [3].

The development of such neural networks and machine learning, in general, would be unimaginable without the invention of the modern computer [48]. Living in the age of laptops, smartphones and smartwatches and constantly observing arrivals of new technologies can give the expectation of endless growing computational potential. The so-called *Moore's law* [49], describes the observation that the number of transistors on a silicon chip doubles every 18 to 24 months. However, the inevitable end of this law was predicted, and indeed a turning point was reached in 2009 when reducing the transistor dimensionality could not improve the device performance any more [50]. The demand for new device architectures and information processing methods arose and is even more motivated through the exponentially increasing amount of data created every day [51].

One promising candidate can be found in the field of *quantum information* [52], the study of information processing tasks that can be carried out using quantum mechanical systems: *quantum computers* containing hundreds of quantum bits became experimentally realisable in the last years [53–55] and give the opportunities to exploit the laws of quantum mechanics to avoid the limits of classical computing.

These quantum bits, called *qubits*, are physically implemented as two-state devices. Since these have different characteristics than their classical counterparts, it is possible to find quantum algorithms which outperform classical computers for specific tasks [56–58]. Today's quantum computers do not yet comprise of enough qubits to run the most useful quantum algorithms. However, the arrival and public access [59] of these devices is a huge motivation for research in many scientific areas.

One of them is *quantum machine learning* (QML) [60–62], whose invention was motivated by the great success of classical ML and the big hope invested in quantum computers. Generally we can divide the field into three categories: classical ML with quantum data [63–65], quantum speed-ups for classical ML [66–69], and quantum algorithms used on quantum data [70–79].

In the first category, quantum data is fed into a classical machine learning algorithm, for instance, to construct representations of many-body systems [65] or estimating physical parameters in quantum metrology [63]. On the contrary, quantum speed-ups for classical ML are motivated with the hope that quantum information processors, which produce patterns which are classically challenging to create, can probably also recognise patterns, which are difficult to distinguish classically [60]. Here, for example, subroutines of an otherwise classical algorithm are quantised [66] using classical data encoded as quantum states.

In this work, we focus on the last category, which raises the most questions. Here both, the algorithm and the training data are based on quantum mechanics, and in analogy to the classical case, *quantum neural networks* (QNNs) [74, 77, 79–95] are a widely spread tool for solving quantum tasks for which no generic quantum algorithm exists. Furthermore, exploiting such QNNs with ML allows to characterise quantum states and operations with fewer samples, compared to

quantum state tomography [96], or quantum process tomography [97], where the number of needed samples increases exponentially with the number of particles [98] and the characterisation of even today’s still minimal quantum devices is impossible. Here QNNs can be explicitly useful since they allow to use quantum devices themselves to cope with large amounts of produced quantum data.

However, the search for the most optimal quantum versions of a neuron, network structure or training algorithm is still ongoing, and the advent of the above mentioned early quantum computers motivates scientists to find executable implementations of QNNs. This work addresses these questions and presents so-called *dissipative quantum neural networks* (DQNNs) [79]. Their training algorithm is designed for fully quantum learning tasks and allows efficient optimisation with memory requirements scaling only with the width, not the length of the QNN. Moreover, we demonstrate that it can be successfully implemented on today’s quantum computers [99].

Due to the rapid progress in quantum learning theory, it is also important to understand the ultimate limits for training methods. Therefore we present the *quantum no free lunch* (QNFL) theorem [100]. It describes a bound on the probability that a quantum device, which can be modelled as a unitary process and is trained with quantum examples, gives an incorrect output for a random input. This theorem gives us a tool to review the prior demonstrated DQNNs.

The DQNNs discussed before are exclusively trained with training data pairs in form of input and desired output states. We extend this ansatz in two directions. Quantum data is always structured due to the structure of the device producing it. In the first extension we present a variation of the DQNN training algorithm using the graph structure of quantum data [101] and demonstrate that we can improve the learning efficiency and the generalisation behaviour by including this additional information.

Whereas both the original DQNN and the DQNN, including graph structure, focus on characterising an unknown quantum operation while learning from the data pairs, we further undertake the task of extending a set of quantum states to states which have similar properties. Hence, in the second extension we follow a generative adversarial approach, where two DQNNs, a generator and a discriminator model, are trained in a competitive manner [102]. Learning characteristics of the training data, the generative model is able to produce quantum states with similar properties. The resulting extended quantum data set can be, for example, useful for training other QNN architectures or experiments.

## Outline

Since the topic is interdisciplinary, we begin with two introductory chapters, whereof one or both can be skipped based on the experience and knowledge of the reader. In **Chapter 2** we introduce classical artificial neurons, neural networks and their training algorithms. This helps to follow the discussion of their quantum analogues. Further, we give an overview of standard methods and applications.

In contrast, **Chapter 3** introduces the field of quantum information. Here we

## 1. Introduction

not only present the characteristics of qubits but also how these can be used in quantum algorithms to outperform classical algorithms. Moreover, we introduce quantum circuits which provides the basis for the implementation of QNNs on today's quantum devices. It is followed by a discussion of quantum computers and an introduction to quantum neural networks.

The heart of this thesis is **Chapter 4**, where the DQNNs are introduced. Here we present an analogue of the classical neuron and explain the architecture of these kinds of QNNs. We define loss functions, which are optimised during the training algorithm. Moreover, we not only show the results of a classical simulation but also of the implementation on actual quantum computers. At the end, we discuss the training behaviour of the DQNN compared to another QNN architecture [103–105]. The following part, **Chapter 5**, contains an introduction to the classical NFL theorem and the derivative and application of the QNFL theorem.

In **Chapter 6**, we give a short overview on how graph-structured data is used in classical ML. Afterwards, we explain and demonstrate how to include the graph structure of quantum training data into the optimisation algorithm of the DQNN. Therefore, we formulate new loss functions and rules to update the network during the training and present their usage in classical simulations and quantum device implementation.

We refer to the generative adversarial NNs in **Chapter 7**. After explaining their basic concepts in classical ML, we transfer the model to QNNs and construct generative adversarial DQNNs. Further, we show that these are able to extend data sets to quantum states which have similar properties, given only a few samples. Here we include classical simulations as well as simulations of quantum devices. Finally, **Chapter 8** concludes our results and gives an outlook of potential further research directions.

# 2

---

## Classical neural networks

---

For all vertebrate and also most invertebrate animals, the brain is the nervous system's centre. This complex organ and also the whole nervous system is built of billions of fundamental units, referred to as *neurons* [106]. The connection of one of these building blocks to others through so-called synapses allows interactions. Whereas the connections within smaller groups of neurons can be recorded, studying the communication between a larger population of these units is very tough [107].

Therefore in the middle of the 20th century, the first computational models for neural networks were proposed [108–110]. Based on these ideas, artificial neural networks, also abbreviated by *neural networks* (NNs) [111–113] arose and could be performed on the at this time available electronic computers. Soon, NNs were applied to describe biological neural networks. Further, the usage for AI got more attention.

Nowadays, NNs built of many layers of neurons are highly used tools for machine learning (ML) [1, 3, 4, 47] and applied in endless different fields of research and sectors of industry [114]. We will discuss some of them throughout the chapter when presenting different NN methods.

In the following we will focus on NN build of simple consecutive layers trained with training data pairs but also refer to other techniques. We start our discussion by explaining the building blocks and different architectures of NNs in Section 2.1 and Section 2.2, in order to explain the training process later on. It follows an introduction to training data and its usage in Section 2.3. We will further discuss different optimisation techniques used for training the NN architectures in Section 2.4. Since this work's limit allows only giving a brief overview into the huge world of NNs, we refer to [1, 3, 4] for a comprehensive discussion.

## 2. Classical neural networks

### 2.1 Artificial neurons

The *perceptron*, developed by Frank Rosenblatt in 1958 [115], was the first artificial neural network with complex adaptive behaviour [113]. Until today versions of this building block, referred to as *artificial neurons*, are used for NN architectures.

Such a neuron takes  $n$  inputs  $\{x_1, \dots, x_n\}$  and has a single binary output  $y$ , also called *activation* as depicted in Figure 2.1. Every input has an assigned weight  $w_i \in \mathbb{R}$ . Additionally, the neuron is equipped with a bias  $b \in \mathbb{R}$ . The neuron's output is computed through

$$y = \kappa(z) = \kappa\left(\sum_i w_i x_i + b\right), \quad (2.1)$$

where  $\kappa(z)$  denotes a so-called *activation function*.




Figure 2.1.: **Artificial neuron.** The building block of NNs takes  $n$  inputs and outputs an activation  $y$  using the activation function  $\kappa$ .

Rosenblatt originally allowed only binary inputs  $x_i \in \{0, 1\}$  and used the *step function*, see Figure 2.2a, as the activation function  $\kappa$  for the definition of the perceptron. Hence, the output is binary. To intuitively understand the method of a neuron, we can imagine the perceptron's task as deciding between two choices, 0 and 1. The inputs  $x_i$  can be seen as arguments with different importances  $w_i$ , where an argument with  $w_i x_i < 0$  is pro-choice 0 and  $w_i x_i > 0$  is pro-choice 1, respectively. The step function and the bias  $b$  describe a threshold. Depending on which side of the threshold the weighted sum  $\sum_i w_i x_i$  of the arguments is, the perceptron “decides” for the output 0 or 1.

#### Activation functions

In the context of NNs, where neurons are layered, we can say that the step function, used in the original definition of the perceptron, *activates* the neuron if the input is above a certain threshold. Otherwise, the neuron is *deactivated*, which means the input will not propagate further through the network, as will be described in Section 2.2. It turns out that using the step function as the activation of a neuron



causes problems: sometimes, a slight change in the weights or bias of any single perceptron in the network can flip the perceptron's binary output and therefore also changes completely the output of the NN. This behaviour can damage the training process [3], and it is aimed that small changes in the input cause only minor changes in the output of NNs. This can be realised by choosing activation functions different from the step function.

We can generally say that choosing the activation functions wisely is crucial for good training results since the activation functions decide if an input of a neuron is relevant. To better understand the activation process, we discuss the advantages and disadvantages of some standard activation functions in the following. An overview of the named functions can be found in Figure 2.2.

For simplicity, we restrict the discussion to activation functions with one input. This input is often chosen to be the weighted sum  $z = \sum_i w_i x_i + b$  of the preceding layer's outputs  $x_i$ . There exist similar functions with many inputs, defined directly on the outputs  $x_i$ . One example is the *softmax* function [4] defined for  $n$  outputs  $z_i$  via

$$\kappa_{\text{softmax}}(z_1, \dots, z_n)_i = \frac{\exp^{z_i}}{\sum_{j=1}^n \exp^{z_j}}.$$

The most straightforward smooth activation would be a *linear function*, for example, the identity. Nevertheless, non-linear activation functions are usually chosen for the activation of neurons, since linear activation functions are problematic. First of all, the derivative of a linear function is a constant. As we will see later when explaining the back-propagation algorithm in Section 2.4, the derivative is used to update the neurons. If the derivative is a constant, it does not include any information about how the network operates on the input data. On the other hand, compositions of linear functions are linear functions again. Thus the whole network, built of many of these neurons, technically collapses to a single-layer network.

On the contrary, non-linear functions create more complex mappings. A common choice is the differentiable and monotonic *Sigmoid function* [116, 117], also named *logistic function*. Whereas it transforms the output smoothly into a number  $\in [0, 1]$ , the mean disadvantage of this function is that its gradient is vanishing for very low or high input values. This stops the network's learning process, as will be clear in Section 2.4. Nevertheless neurons using the Sigmoid function for activation, also commonly called *Sigmoid neurons*, find applications, for example in classification tasks [118, 119].


The *hyperbolic tangens* (TanH) is similar to the Sigmoid function, but often preferred over it, because it is symmetric around the origin, which makes its output values centred around the zero value. One can express the hyperbolic tangens using the Sigmoid function, namely

$$\kappa_{\text{TanH}}(z) = \frac{2}{1 + e^{-2z}} - 1 = 2\kappa_{\text{sigmoid}}(2z) - 1.$$

Hence the problems with vanishing gradients are the same.


## 2. Classical neural networks

$$\kappa_{\text{step}}(z) = \begin{cases} 0 & \text{if } z \geq 0 \\ 1 & \text{if } z < 0 \end{cases}$$




(a) Binary Step

$$\kappa_{\text{lin}}(z) = \alpha z$$




(b) Linear function

$$\kappa_{\text{sigmoid}}(z) = \frac{1}{1 + e^{-z}}.$$




(c) Sigmoid/logistic function

$$\kappa_{\text{TanH}}(z) = \frac{2}{1 + e^{-2z}} - 1$$




(d) Hyperbolic tangens

$$\kappa_{\text{ReLU}}(z) = \max(0, z).$$



(e) Rectified Linear Unit

$$\kappa_{\text{LReLU}}(z) = \max(\alpha z, z).$$



(f) Leaky Rectified Linear Unit

Figure 2.2.: **Activation functions.** This figure presents the graphical and mathematical representation of a selections of activation functions.

The *rectified linear unit* (ReLU) deactivates the neuron if the output is smaller than zero. Otherwise, a linear function, for example, the identity, acts. This activation function is computationally very inexpensive. Since the gradient becomes zero for non-positive input values, training algorithms based on the gradients fail in these cases.

The *leaky rectified linear unit* solves this problem, with slightly descending outputs for inputs smaller than zero, see Figure 2.2f. A further variation of this activation function is the *parametrised* ReLU, where the parameter  $\alpha$  is optimised during network training. All versions of ReLU have the disadvantage that the function's output is not always in  $[0, 1]$ , which is often desirable.

## 2.2 Network architecture

After the first enthusiasm about perceptrons, it got rapidly clear that these one-layer NNs were quite limited in computational power [120]. It was discovered that stacking these early artificial neurons in layers increases the computational power. Whereas with only one of the by Rosenblatt introduced perceptrons, only the learning of linearly separable classes can be performed, i.e. a linear hyperplane can divide two classes, with *multi-layer perceptrons* also non-linear classification problems can be solved [121, 122].

In this section, we describe how to build such multi-layer NNs out of neurons. Since we will only mention a few different NN architectures in the following, but many more exist, we point to [123] for a more complete overview.




Figure 2.3.: **Feed-forward NN.** We depict neurons as circles. This figure shows a feed-forward NN with  $L$  hidden layers. The connections between the neurons are denoted by arrows.

## 2. Classical neural networks

Most often, the neurons in NN are arranged layerwise, see Figure 2.3. We call the first layer of neurons, which get the initial input, the *input layer*. The last layer of neurons is named *output layer*. This layer's output is the final output of the network. The layers in between are called *hidden layers*.

Further, we call networks with many hidden layers *deep NNs* [4]. While working through a considerable amount of the layers, the original information gets more and more abstract, and in that way, the complex data processing is divided into a series of simple nested assignments.

### Feed-forward neural networks

The simplest NN architecture can be found in *feed-forward neural networks* [124]. Here the neurons get the output of previous layers neurons as an input, and no loops are built-in. Furthermore, such NNs are often built of *fully connected layers*, i.e. layers where all the inputs from one layer are connected to every neuron of the next layer. See Figure 2.3 for a depiction.

Despite this simple structure, these networks are widely used, for example, for modelling the spread of COVID-19 [125], forecasting of wind power [126], or studying the drying kinetics of pistachio nuts [127]. Since many applications of feed-forward NNs are based on pattern recognition [128–130], we will describe how handwritten digits can be classified using a supervised training ansatz and a feed-forward neural network in Section 2.3.

### Recurrent neural networks

*Recurrent neural networks* (RNNs) [4, 47, 131–133] are constructed in a more complicated way than the feed-forward NNs we discussed above. Such RNNs can be used for working with data sequences, which can be, for example of temporal order as in speech recognition [134–136], video analysis [137–139] or language processing tasks [140, 141]. Instead of, for instance, the whole text is being fed into the algorithm as one single input, the RNN allows several inputs  $x_t$  during the time the algorithm is running, for example, the single words, and gives also outputs  $y_t$  during the running process. These outputs are not only based on the according input but also on the hidden layers of the preceding time step  $t - 1$ . Figure 2.4 shows the simplest version of an RNN.

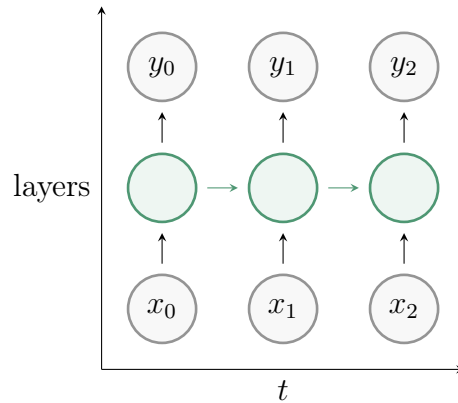


Figure 2.4.: **Recurrent neural network.** This figure represents an RNN consisting of a one-neuron input, hidden (colored) and output layer in three steps  $t \in \{0, 1, 2\}$ .

### Convolutional neural networks

Yet another NN structure can be found in *convolutional neural networks* (CNN) [4, 47, 142–144]. Besides fully connected layers, as we used in the feed-forward NNs, also so-called *convolutional layers* are used to shape networks of this class. These layers use the convolution of the layer’s input, often inserted in matrix form, with another matrix, often called *kernel*.

A common example is the input of a grey-scale image in form of a matrix. Sharpening the image can be executed by executing a convolution between the image matrix and a kernel, which is in this case a sharpening matrix, see Figure 2.5.

3	3	4	3	3	3	sharpening kernel	-1	-1	-1	*	-23	-37	-42	-38	
3	4	5	6	4	4		-1	5	-1		=	-49	-24	-30	-24
3	4	18	20	15	4		-1	-1	-1		-44	-5	78	-47	
4	7	25	43	14	4										
3	4	18	16	11	4										
3	4	4	6	9	4										

Figure 2.5.: **Matrix convolution.** In this example for a convolution of a matrix with a sharpening kernel the computation of one matrix element is highlighted for demonstration. The number 78 is the sum of the element-wise multiplications of the highlighted matrix elements.

In many applications, additional layers which reduce the number of parameters of their input are used. These layers are referred to as *pooling layers*. An example for such a filter is *max pooling*, where the filter selects the biggest value of an array, exemplary depicted in Figure 2.6. Furthermore, *average pooling* is common, where every array is replaced by its average value.

## 2. Classical neural networks




Figure 2.6.: **Max pooling.** An example for pooling is max pooling, where the filter selects the biggest element of every array.

A typical CNN starts with a convolutional layer, followed by additional convolutional layers or pooling layers and ends with a fully connected output layer. The architecture of CNNs was inspired by the neurobiological model of the visual cortex [145].

Since CNNs are good tools for analysing image data [146], nowadays CNNs have a big impact on many domains, for example, health informatics [31]. These kinds of NNs are used, for example, on tomography-computed images to categorise them by body-parts [33,34] or to classify diseases [35]. Other applications are based on X-ray images, for instance to classify tuberculosis [36] or diagnosis of COVID-19 [37,38]. Furthermore, magnetic resonance images of the brain are processed [39,40].




Figure 2.7.: **GAN.** The generative NN  $G$  produces data, whereas the discriminative NN  $D$  has the training goal to distinct between the by  $G$  produced data and the training data.

### Adversarial neural networks

The task of handwritten digit recognition was mentioned above and will be studied in the following two sections. Such a problem is a typical classification task. A different goal would be to train a NN to be able to *produce* such digit images. We call such a NN model *generative*. An adversarial process, where two NNs are trained, see Figure 2.7, can be used to train such a *generative adversarial network* (GAN) [147]: the generative model  $G$  captures the data distribution and produces data, whereas the other model, referred to as *discriminative model*  $D$  estimates

the probability  $D(\vec{x})$  that a sample  $\vec{x}$  came from the training data rather than from  $G$ . Both of the NNs have different training goals and are well trained at the end. Therefore, we are left with a perfectly well-trained generator. We will explain the training process in detail in Section 7.1.

## 2.3 Training data and loss functions

So far, we only described the information flow of input data through the NN. The neurons, and hence the NN, are parametrised through weights and biases, see Equation (2.1), and sometimes parameters of the activation function, for example, like in the parametrised ReLu, defined in Figure 2.2f. What is missing now is an algorithm to optimise the parameters in a way that the NN satisfies our standards.

To explain the comparison on which the parameter's update is based, we first have to define a training goal. As mentioned in the introduction, we focus on *supervised learning* tasks in this chapter. These kinds of learning tasks are based on *training data*, most often structured in pairs, each containing an input and the desired output value.

At this point, an example comes in handy: one of the most, if not *the* most famous data set is the MNIST data set introduced in [146] and available at [148]. This set contains thousands of greyscale  $28 \times 28 = 784$ -pixel images of handwritten digits and the suiting label of the set  $\{0, 1, \dots, 9\}$ .

Based on this data set, the task, a typical classification problem, is clear: we want to train the NN so that we can feed it an unseen image saved as a vector  $x$  of length 784 and want to get the correct digit of the set  $\{0, 1, \dots, 9\}$ . Due to the characteristics of the training algorithm, we convert the digits into vectors of dimension 10. See Figure 2.8 for an example.

$$x_i \hat{=} \mathbf{2}, \quad d_i = (0, 0, 1, 0, 0, 0, 0, 0, 0, 0)^T$$

Figure 2.8.: **MNIST data set.** In this figure an exemplary input, a 784-pixel image, and desired output, a vector in  $\mathbb{R}^{10}$  are depicted.

To test the training process it makes sense to divide the provided training data set of length  $N$  into a *supervised training set*  $\mathcal{S}_{SV}$  containing  $S$  pairs and in a *unsupervised validation set*  $\mathcal{S}_{USV}$  build of the remaining data. The latter stays unseen by the network during the training and is only used for validation. We can denote the two sets as

$$\begin{aligned} \mathcal{S}_{SV} &= \{\{x_i, d_i\}\}_{i=1}^S \\ \mathcal{S}_{USV} &= \{\{x_i, d_i\}\}_{i=S+1}^N. \end{aligned}$$

The original MNIST data contains, for example,  $S = 60000$  training pairs and  $N - S = 10000$  testing pairs [148].

## 2. Classical neural networks

### Training and validation loss

Given a NN suiting the MNIST example, i.e. one with a 784-neuron input layer and a 10-neuron output layer, we can load an image saved as  $x_i$  to the network and feed-forward the information as described in Section 2.2. We suppose further that we get a vector of dimension 10 as an output, consisting of the by the NN estimated probabilities that  $x_i$  has a specific label. We denote the process with a function  $\mathcal{E}_{w,b} : \mathbb{R}^{784} \rightarrow \mathbb{R}^{10}$ , which depends on the weights and biases, saved in vectors  $w$  and  $b$ . The output of the network can then be denoted as  $\mathcal{E}_{w,b}(x_i)$ .

Independently from the supervised training task it is indispensable to compare the network's output  $\mathcal{E}_{w,b}(x_i) \in \mathbb{R}^{10}$  with the desired output  $d_i \in \mathbb{R}^{10}$ . This is done by a *loss function*, also referred to as *cost function* or *objective function*. Not only should this function compare the two pieces of information, but it should also converge to a global extreme point, which is achieved if the training goal is accomplished. Further, it is optimal if the gradient of the loss is continuous since the below-presented algorithm is based on it.

A common loss function is the *mean squared error* (MSE) [149, 150], also often simply called the *quadratic cost*. Given the network's output  $\mathcal{E}_{w,b}(x_i)$  and the desired output  $y_i$  we can use the MSE to form the *training loss*

$$\mathcal{L}_{\text{SV,MSE}}(w, b) = \frac{1}{S} \sum_{i=1}^S \|d_i - \mathcal{E}_{w,b}(x_i)\|^2. \quad (2.2)$$

If this loss functions gets small during training we are sure that the algorithm found weights and biases which suite the aim. For validation, we can define the *validation loss*

$$\mathcal{L}_{\text{USV,MSE}}(w, b) = \frac{1}{N - S} \sum_{i=S+1}^N \|d_i - \mathcal{E}_{w,b}(x_i)\|^2. \quad (2.3)$$

Note that the validation is a not neglectable part of the training since, in some cases, the training leads to parameters that let the NN perform very well on the training data set but cannot generalise the same performance to unseen data. We refer to this issue as *overfitting* and can detect it with the validation data set. Overfitting can be avoided with *early stopping* [4], where the implementation includes stopping the training to the time where the validation loss  $\mathcal{L}_{\text{USV}}$  begins to rise again, although the training loss error  $\mathcal{L}_{\text{SV}}$  still gets smaller.

### Choice of the loss function

The choice of the loss depends on the learning problem and method. For example, the MSE is very sensitive towards outliers, due to the quadratic behaviour. In cases where this causes problems, for example where the input is not normally



### 2.3. Training data and loss functions

distributed, the *mean absolute error* (MAE), namely

$$\mathcal{L}_{\text{SV,MAE}}(w, b) = \frac{1}{S} \sum_{i=1}^S |d_i - \mathcal{E}_{w,b}(x_i)|,$$

can be an alternative. Note that the validation loss can be defined for MAE in analogy to Equation (2.3). This also holds for the loss functions presented in the following.

Another common loss function for regression problems is the *smooth mean absolute error*, also called *Huber loss* [151]. This loss function is defined piecewise and depends on an additional parameter  $\zeta$ , namely

$$\mathcal{L}_{\text{SV,Huber}}(w, b) = \frac{1}{S} \sum_{i=1}^S \begin{cases} \frac{1}{2}(d_i - \mathcal{E}_{w,b}(x_i))^2 & \text{for } |d_i - \mathcal{E}_{w,b}(x_i)| \leq \zeta \\ \zeta(|d_i - \mathcal{E}_{w,b}(x_i)| - \frac{1}{2}\zeta) & \text{otherwise.} \end{cases}$$

Loss functions containing the entropy were proved to be useful for classification tasks [152–155]. The most simple case is binary classification, where we have training data pairs, containing an input  $x_i$  and a label  $d(x_i) = 0$  or  $d(x_i) = 1$ . The network's output for such an input is the probability  $\mathcal{E}_{w,b}(x_i) = p_{x_i}(0)$  that the input is labelled with  $d(x_i) = 0$ . Since  $1 - p_{x_i}(0) = p_{x_i}(1)$  we can use the *binary cross entropy* (BCE) as loss function, i.e.

$$\mathcal{L}_{\text{SV,BCE}}(w, b) = -\frac{1}{S} \sum_{i=1}^S d(x_i) \log(\mathcal{E}_{w,b}(x_i)) + (1 - d(x_i)) \log(1 - \mathcal{E}_{w,b}(x_i)).$$

From this definition, it gets clear that minimising this loss function leads to a correctly labelling NN.

When working with more than two, namely  $C$ , different classes, the class labels are unit vectors, similar to the example in Figure 2.8. The NN's output has to be a  $C$  dimensional vector  $\mathcal{E}_{w,b}(x_i) = p_{x_i}$  with  $p_{x_i,j} \in \mathbb{R}_0^+$ . Here the general *cross entropy* (CE) is used to formulate the loss function

$$\mathcal{L}_{\text{SV,CE}}(w, b) = -\frac{1}{S} \sum_{i=1}^S \sum_{j=1}^C d(x_i)_j \log(\mathcal{E}_{w,b}(x_i)_j).$$

As an example we can assume  $C = 3$  classes and the label of a specific training data  $i$  input is  $d(x_i) = \{0, 1, 0\}$ , i.e. it belongs to the second class. The network predicts the probabilities  $p_{x_i} = \{0.2, 0.5, 0.3\}$ , hence 50 percent for the correct label. The according summand in  $\mathcal{L}_{\text{SV,BCE}}$  is then  $0 \times 0.2 + 1 \times 0.5 + 0 \times 0.3$ .

Since classification tasks are a vast field in ML, many other loss functions were studied. Alternatives to the CE are, for example, the *pairwise loss* [156], which runs over pairs of training data inputs, or the *triplet loss* [157] working with the assumption that a training data input is closer to all inputs with the same label, then to those which are labelled differently. More definitions of loss functions can

## 2. Classical neural networks

be found, for example, at [4].

### Unsupervised learning approaches

The so far discussed approaches are supervised [158–160]. These model the relationships or patterns between the input features and the target prediction output. At the end of the process, the NN can imitate the relation and predict output values for unseen data. For these kinds of models, having enough and proper training data is key.

Beyond that, also *unsupervised* algorithms [161–163] are used. These methods do not rely on labelled input and output data and are especially useful in cases where it is unclear what kind of patterns to look for in the data. An example of unsupervised learning is *k-means clustering* [164–166]. The algorithm aims to partition data into  $k$  clusters given a set of vectors and minimises the within-cluster sum of squares. Note that  $k$  is fixed beforehand. One method used for finding the number of clusters is, for example, the *elbow method* [167].

There are also hybrid versions of unsupervised and supervised learning. These are often used for classification problems, where only a small subset of the data is labelled because obtaining labels is expensive or impossible. In these algorithms, the classification is based not only on the labelled part of the data but also on all data bits. Applications for *semi-supervised learning* [168, 169] are for example in speech analysis [170], image search [171] or genomics [172].

An examples for a semi-supervised ansatz will be presented in Section 6.2, where *machine learning with graphs* [173–175] is discussed. A graph describes the connection between different vertices, and only some of the vertices are labelled. These labels can be, for example, of the form of input and output data pairs as we met them in the above explained supervised ansatz. In addition to the supervised data, the training algorithms then also use the information implemented via a graph structure.

A third large area of machine learning alongside supervised and unsupervised learning is *reinforcement learning* [176]. This method is modelled as a *Markov decision process*, where we have a set of environment and agent states, a set of actions and the probability of transition from one state to another under a specific action. Further, some rules describe what the agent observes, and a (positive or negative) reward is given after transitioning from one state to another with a specific action. The training goal is to maximise the sum of rewards.

## 2.4 Optimisation

We have already argued that layering neurons is very beneficial in contrast to training single neurons. However, these multi-layer networks were not helpful first since no suiting training algorithm could be found. This changed with the invention of the *back-propagation algorithm* [177–179], which will be discussed at the end of this section.

The all-over optimisation procedure using such an algorithm can be defined as follows: we feed information through the network and get output data. This output data is compared with some desired output, as discussed in Section 2.3. Based on this comparison, the parameters are updated. This process is repeated until the desired accuracy is reached. We will describe the update of the parameters in the following, assuming to train a simple feed-forward NN.

### Gradient descent

The training algorithm improves the network by changing the before training randomly initialised parameters  $w_0$  and  $b_0$ . By definition this is done by optimising the training loss discussed in Section 2.3, i.e. finding a global extreme point of  $\mathcal{L}_{\text{SV}}(w, b)$ .

The most common method for this purpose is *gradient descent*, which was proposed for non-linear optimisation problems by [180] based on earlier contributions of Augustin-Louis Cauchy and Jacques Hadamard. With this technique, a local minimum can be found by changing the parameters  $w_t$  and  $b_t$  into the direction of the gradient of the loss function concerning the parameters, namely

$$\begin{aligned} w_{t+1} &= w_t - \eta \nabla_{w_t} \mathcal{L}_{\text{SV}}(w_t, b_t) \\ b_{t+1} &= b_t - \eta \nabla_{b_t} \mathcal{L}_{\text{SV}}(w_t, b_t), \end{aligned}$$

where  $t$  denotes the training *step* or *epoch* and  $\eta$  the *learning rate*. Although there is mathematically no guarantee for success, numerical experiments have shown that the method reliably finds global minima of deep NNs [181].

### Alternatives to gradient descent

Finally, we want to name some optimisation methods different from the above-described gradient descent method. A more detailed overview of such methods and their variations can be found at [4], the comparison of different proposals in [182, 183]. We can summarise the gradient descent method via the rule

$$\vartheta_{t+1} = \vartheta_t - \eta \nabla_{\vartheta_t} \mathcal{L}_{\text{SV}}(\vartheta_t),$$

for updating a parameter  $\vartheta$  and name some adaptations of this rule in the following.

To optimise the computing time of each training step, *stochastic gradient descent* (SGD) can be utilised. Whereas the in Section 2.4 described method uses the loss function evaluated on all  $S$  training pairs, here for each training epoch only one random training pair is chosen, the loss function evaluated and the gradients computed [184]. A good compromise is *mini-batch stochastic gradient descent*, where a randomly chosen set of  $M < S$  training pairs is selected and used instead of just one [185].

A faster converging technique is the *momentum method* [186, 187], where we update the parameter is done via  $\vartheta_{t+1} = \vartheta_t + M(t)$  while preserving the momentum

## 2. Classical neural networks

of the last update via an additional parameter  $\mu$  by using

$$M(t+1) = \mu M(t) - \eta \nabla_{\vartheta_t} \mathcal{L}_{\text{SV}}(\vartheta_t).$$

For very steep parts of the parameter change, the momentum gets bigger and finding the extrema is accelerated. On the other hand, with this method, the change of the parameter gets small if we get near the turning point, and overshooting it is less likely.

So far, we treated the learning rate  $\eta$  as a beforehand fixed and unchanged parameter during the training. On the contrary, some optimisation methods use *adaptive learning rates* changing individually for each parameter during the training [4]. One of the first proposed adaptive learning technique was the *delta-bar-delta algorithm* [188], which is based on a simple idea: if the partial derivative of the loss with respect to the parameter remains of the same sign, the learning rate increases, otherwise it decreases.

A very common optimisation algorithm using adaptive learning rates is *Adam* [189], a synonym for *adaptive moments*. It can be seen as an improvement of two previously invented methods, *AdaGrad* [190], and *RMSProp* [191]: Using AdaGrad, the learning rates scale inversely proportional to the square root of the sum of all squared values of the earlier epoch's gradients. RMSProp is a slight improvement of this technique as it uses the squared gradients of earlier epochs primarily. The method Adam is based on RMSProp, but also includes techniques of the momentum method.

### Back-propagation algorithm

When it comes to calculating the gradients of the loss needed for the parameter update, it is often referred to the already above-mentioned *back-propagation algorithm* [177–179]. This algorithm can be applied for loss functions that can be written as averages over the loss for a specific training pair since the derivatives of the training loss are computed through the averaging of the derivatives of the training loss of single training examples. Further, the loss has to be a function of the NN's output  $\mathcal{E}_{w,b}(x_i)$ . Both is the case for the in Equation (2.2) defined training loss. A very detailed and graphic description can be found in [3]. In the following, we will present a summary.

Since the parameters denoted in the vectors  $w$  and  $b$  change during the training epochs, their value depends on  $t$ . Due to that also the training loss  $\mathcal{L}_{\text{SV}}$  and the activations  $y^l$  depend on  $t$ . Note that for convenience, we drop the parameter  $t$  in the description of the gradient evaluation for a specific epoch, in the following.

For every layer  $l$  we save the weights of the neurons connection in a matrix  $w^l$ , where  $w_{j,k}^l$  is the weight from the  $k$ th neuron in layer number  $l-1$  to the  $j$ th neuron in layer  $l$ . See Figure 2.9 for an example. Further, we describe the biases in in vectors  $b^l$ , where  $b_j^l$  describes the bias of the  $j$ th neuron in the  $l$ th layer.

Since we work with feed-forward NNs, we can use  $w^l$ ,  $b^l$  and the activation function  $\kappa$  to express the activation of layer  $l$  through the activation of layer  $(l-1)$ ,




Figure 2.9.: **Notation for back-propagation algorithm.** Two layers  $l - 1$  and  $l$  of a feed-forward NN containing two and three neurons are connected via weights  $\{w_{11}^l, \dots, w_{23}^l\}$ .

namely

$$y^l = \kappa(z^l) = \kappa(w^l y^{l-1} + b^l),$$

where the component  $z_j^l$  of  $z^l$  is the weighted input to the activation for neuron  $j$  in layer  $l$ .

Since the loss function describes the difference between the desired output and the networks output we define the *error* of the  $j$ th neuron in layer  $l$  as

$$\delta_j^l = \frac{\partial \mathcal{L}_{\text{SV}}(w, b)}{\partial z_j^l},$$

and describe the errors in layer  $l$  as vector  $\delta^l$ .

In the following we will use  $a \odot b$  to denote the element-wise product, i.e.  $(a \odot b)_i = a_i b_i$ , also referred to as Hadamard product. Due to the chain rule we can express the error of the last layer, namely layer  $L$ , as

$$\delta^L = \nabla_y \mathcal{L}_{\text{SV}}(w, b) \odot \frac{\partial \kappa(z^L)}{\partial z^L}.$$

As the name of the algorithm already revealed we can now back-propagate this error through the NN and hence compute the error of an arbitrary layer  $l$  using the error of the preceding layer  $l + 1$ , in other words

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \frac{\partial \kappa(z^l)}{\partial z^l}.$$

## 2. Classical neural networks

The components of the gradients  $\nabla_b \mathcal{L}_{\text{SV}}(w, b)$  and  $\nabla_w \mathcal{L}_{\text{SV}}(w, b)$  are thus

$$\begin{aligned} \frac{\partial \mathcal{L}_{\text{SV}}(w, b)}{\partial b_j^l} &= \delta_j^l \\ \frac{\partial \mathcal{L}_{\text{SV}}(w, b)}{\partial w_{jk}^l} &= y_k^{l-1} \delta_j^l, \end{aligned}$$

where  $y^{l-1} = \kappa(z^{l-1})$  is the activation of layer  $l-1$ . An implementation in Python of the back-propagation algorithm used for classifying handwritten digits can be found at [3].

In Chapter 4 we will introduce a quantum version of a feed-forward NN, which can be optimised with an algorithm similar to the classical back-propagation method we described above. Before digging deeper into quantum NNs, we will give an introduction to quantum information in the following chapter.

# 3

---

## Quantum information

---

As described in Chapter 1, in the past decades, the power of classical computers has grown exponentially. Nevertheless, boundaries for classical computers exist [50]. For instance, the description of more sizeable quantum systems is still impossible on today's classical supercomputers since the dimension of quantum systems scales exponentially with the number of basic building blocks.

Consequently, in the eighties, Richard Feynman proposed the concept of quantum computers [192] with the aim to simulate many-body quantum systems with the use of a quantum system of the same size instead of a classical computer. Building on this idea, the field of quantum computing evolved in the last decades and is now one of the most rapidly developing research areas. Until today many quantum algorithms were proposed with the aim of solving a particular task more efficiently than classical supercomputers. The first breakthroughs happened in the nineties, among them the Deutsch-Jozsa algorithm [56], Shor's factoring algorithm [57] and Grover's search algorithm [58].

Whereas the key developments remained on the more theoretical side first, in the last years processing quantum information on quantum computers has become experimentally possible [53–55]. These first quantum computers are called *noisy intermediate-scale quantum* (NISQ) devices, comprise up to a few hundred quantum bits and give the opportunities to test quantum algorithms for their behaviour under high noise levels.

The arrival of these new devices and also their public access [59] lead to many attempts of solving problems on these devices. It should be underlined at this point that computations on quantum computers containing only about 100 qubits can be, in many cases, simulated classically. The emergence of NISQ devices should not be underrated but also not overrated, or as John Preskill appropriately

### 3. Quantum information

formulated in [53]: “The 100-qubit quantum computer will not change the world right away — we should regard it as a significant step toward the more powerful quantum technologies of the future.”

Altogether, we can conclude that we are still at the beginning of the quantum information era. However, instead of feeling disappointed, this should be seen as an even more significant motivation for study and research in the area of quantum information these days, in the view of many so far undiscovered opportunities.

To make the research presented from Chapter 4 onwards accessible to people new to quantum information, we will discuss the basic definitions in this chapter. We start by introducing the quantum bit and its characteristics differing from the classical bits in Section 3.1. Based on this, we define quantum gates and circuits in Section 3.2. In Section 3.3 we not only deliver insight into why quantum algorithms can outperform classical algorithms but also give an overview of the most famous quantum algorithms. Section 3.4 explains how quantum circuits can be implemented and describes the state of the art of quantum computers. Since the topic of this thesis is the study of quantum neural networks, we conclude with a review of these in Section 3.5. For a wholesome introduction to quantum information, we point to [52].

## 3.1 Quantum bits

In classical information theory, basic binary units of information, called bits, are used for information processing. These bits - physically implemented with a two-state device - can only have one of two values, commonly represented as either 0 or 1. The quantum equivalent is named *qubit*. It represents a two-level quantum mechanical system, often denoted in terms of  $|0\rangle$  and  $|1\rangle$ . These levels can, for example, be two different polarisations of a photon (vertical and horizontal polarisation), two different alignments of spin (up and down) or two states of an electron orbiting an atom (ground state and excited state). But, they can also be defined by more complex systems based on very cold superconducting electrical circuits in which several electrons move [52, 53].

### Superposition

Whereas the state of a classical bit can be either 0 or 1, the qubit is allowed to be in a coherent *superposition* of both levels  $|0\rangle$  and  $|1\rangle$  simultaneously. For a mathematical description of these qubit *states* we identify the levels with two orthonormal basis vectors

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (3.1)$$

A qubit in a so called *pure state* can then be described as a normalised vector  $|\phi\rangle$  in a two-dimensional complex Hilbert space  $\mathcal{H} = \mathbb{C}^2$ . Using the in Equation (3.1)



defined basis states we can express such a state as superposition

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (3.2)$$

where  $\alpha$  and  $\beta$  are complex numbers and  $|\alpha|^2 + |\beta|^2 = 1$ .

In quantum mechanics the *Dirac notation*, read as *bra*  $\langle *|$  and *ket*  $|* \rangle$ , is the standard notation for states. Here  $\langle \psi|$  denotes the conjugate transpose of the vector  $|\psi\rangle$ . The eponymous *bra-ket*  $\langle \psi_1|\psi_2\rangle$  expresses the scalar product of the vectors  $|\psi_1\rangle$  and  $|\psi_2\rangle$ . At this point it is worth to accentuate that this description only includes pure states. Density matrices, discussed later in this section, provide a more general notion of statistical ensembles of pure states.

Since the description of a pure qubit state is a vector with norm 1 we can depict such a state on the surface of a sphere. For this purpose we rewrite Equation (3.2) as

$$|\psi\rangle = e^{i\gamma} \left( \cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\varphi} \sin\left(\frac{\theta}{2}\right) |1\rangle \right). \quad (3.3)$$

The visualisation of quantum states with two numbers  $\theta$  and  $\phi$ , depicted in Figure 3.1, is named *Bloch sphere*. Note that the condition  $|\alpha|^2 + |\beta|^2 = 1$  defines a 3-dimensional sphere. The *Hopf fibration* [193], mapping from the unit 3-sphere to the two-dimensional states space  $\mathbb{C}^2$ , is used to describe the sphere with two parameters  $\theta$  and  $\phi$ , namely as

$$|\psi\rangle_{\text{Bloch}} = \cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\varphi} \sin\left(\frac{\theta}{2}\right) |1\rangle. \quad (3.4)$$

The factor  $e^{i\gamma}$  can be ignored, since it has no *observable effects* which is in detail explained in [52]. We will discuss the abandoning of this total phase later in this chapter again.

Like for the classical bits there are two possible outcomes for the measurement of a qubit, also often taken to be 0 and 1. After measuring in the basis denoted in Equation (3.1) such a qubit state collapses in the pure state  $|0\rangle$  or  $|1\rangle$ , respectively. It is possible to compute the probabilities of getting a special measurement outcome. With probability

$$p_0 = \langle \psi|0\rangle \langle 0|\psi\rangle = |\langle 0|(\alpha |0\rangle + \beta |1\rangle)|^2 = |\alpha \langle 0|0\rangle|^2 = |\alpha|^2$$

the measured value will be 0. We call  $M_0 = |0\rangle \langle 0|$  a *measurement operator*. Analogous we get  $p_1 = |\beta|^2$  with the operator  $M_1 = |1\rangle \langle 1|$ .

To summarise, a qubit can be in an infinite number of possible states, but after the measurement, we again only receive binary information as with the classic bit and the superposition information seems lost. However, with quantum algorithms, this fundamental property of quantum mechanics can be used effectively and is the reason for quantum speed up, as will be explained in Section 3.3.

Although not relevant for this work, we want to mention at this point the existence of the generalised concept of *qudits*. These describe quantum systems

### 3. Quantum information




Figure 3.1.: **Bloch sphere.** Pure qubit states can be represented as points on the surface of a sphere.

with  $d > 2$  levels. In the case of  $d = 3$  the name *qutrit* is established. Albeit most of the studies in quantum information theory operate with qubits, these higher-level bits can be still beneficial in some contexts [194–198].

#### Multiple qubits

In the same way, as classical bits unfold their full potential only when used in strings of many bits, a quantum system and, therefore, quantum computing becomes much more attractive when involving more than one qubit.

We describe  $n$  qubits with the tensor product of  $2^n$  basic states living in the Hilbert space  $\mathcal{H} = (\mathbb{C}^2)^{\otimes n} = \mathbb{C}^{2^n}$ , namely

$$|q_1\rangle \otimes |q_2\rangle \otimes \dots \otimes |q_n\rangle = |q_1 q_2 \dots q_n\rangle, \quad (3.5)$$

where  $q_i$  indicates the state of the  $i$ th qubit with  $q_i \in \{0, 1\}$ .

In that way the basic states of a two qubit system can be denoted as  $|00\rangle$ ,  $|01\rangle$ ,  $|10\rangle$ , and  $|11\rangle$ . A general pure two-qubit state can be expressed as a superposition

$$|\psi\rangle = \alpha_{00} |00\rangle + \alpha_{01} |01\rangle + \alpha_{10} |10\rangle + \alpha_{11} |11\rangle$$

in  $\mathcal{H} = \mathbb{C}^4$ , where  $|\alpha_{00}|^2 + |\alpha_{01}|^2 + |\alpha_{10}|^2 + |\alpha_{11}|^2 = 1$  and  $\alpha_{ij} \in \mathbb{C}$ .

In the same manner a pure  $n$  qubit state can be described as the superposition

$$|\psi\rangle = \sum_{q_1, \dots, q_n \in \{0, 1\}} \alpha_{q_1, \dots, q_n} |q_1 \dots q_n\rangle$$

in  $\mathcal{H} = \mathbb{C}^{2^n}$  with complex coefficients  $\alpha_{q_1, \dots, q_n}$ .

At this point, the exponential growth of the Hilbert space with the number of qubits becomes obvious, and with this, the high costs when simulating quantum systems classically fall into place. On the other hand, this exponential relationship indicates the extensive capacity of quantum computing.

### Mixed states

To this point all mentioned quantum states were pure states, which can be represented as vectors in  $\mathbb{C}^{2^n}$ . Though in general a system has to be described by a mixture of these pure states. Such a *mixed state* can be described with a matrix  $\rho \in \mathcal{H} = \mathbb{C}^{2^n} \times \mathbb{C}^{2^n}$  with positive eigenvalues,  $\text{tr}(\rho) = 1$ , a so called *density matrix*.

If the state  $\rho$  is a mixture of  $m$  pure states  $|\psi_i\rangle$ , we can describe as a convex combination of these states, namely as

$$\rho = \sum_i^m p_i |\psi_i\rangle \langle \psi_i|$$

with  $\sum_i p_i = 1$ . We can view a pure state rather as special case of the mixed state and can write it as

$$\rho = |\psi\rangle \langle \psi|.$$

Note that at this point the total phase of  $|\psi\rangle$  is lost (compare Equation (3.3) and Equation (3.4)). For the density matrix of a pure state it is  $\rho = \rho^2$  and  $\text{tr}(\rho^2) = 1$ .

As explained when using the bracket notation in Equation (3.5), we use the tensor product to build the composite of two quantum systems, for example, two qubits. When working with mixed states, for example, a qubit  $A$  in the state  $\rho_A \in \mathcal{H}_A$  and qubit  $B$  in the state  $\rho_B \in \mathcal{H}_B$ , we can describe the composite of both qubits with the state

$$\rho_{AB} = \rho_A \otimes \rho_B$$

in the Hilbert space  $\mathcal{H}_A \otimes \mathcal{H}_B$ . Not only the composition of systems is frequently used in quantum algorithms, but also the reduction to a subsystem is an important tool: tracing over the space  $\mathcal{H}_B$  gives

$$\rho_A = \text{tr}_B(\rho_{AB}) \in \mathcal{H}_A.$$

The most simple composition of  $n$  systems can be written as

$$\rho_{1,\dots,n} = \rho_1 \otimes \dots \otimes \rho_n.$$

When tracing out the  $i$ th system we get

$$\text{tr}_i(\rho_{1,\dots,n}) = \rho_1 \otimes \dots \otimes \rho_{i-1} \otimes \rho_{i+1} \otimes \dots \otimes \rho_n.$$

We want to close the description of mixed states with two remarks. Since with the density matrix we generalised the description of a quantum state, we

### 3. Quantum information

have to update the way to calculate probabilities: a mixed state  $\rho$  collapses into the state  $|\psi\rangle$  after measurement with probability  $p = \text{tr}(|\psi\rangle\langle\psi|\rho) = \langle\psi|\rho|\psi\rangle$ . More generally we can describe a measurement with an linear operator  $M$ , called *observable*. The expectation value of the measurement outcome is given by  $\text{tr}(\rho M)$  when measuring  $\rho$ . We point to [52] for a complete discussion of observables.

The second remark concerns the representation of one-qubit states in the Bloch sphere, see Figure 3.1. Mixed one-qubit states can be represented in this way as well. Whereas the pure states lie on the surface of the Bloch sphere, a general mixed state lies inside. Hence the representation requires the radius  $r$  and is given by

$$\rho = \frac{\mathbb{1} + \vec{r}\vec{\sigma}}{2},$$

where  $\vec{r} \in \mathbb{R}^2$  is called the Bloch vector and  $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)^T$  is built from the Pauli matrices, which are denoted in Figure 3.2.

## Entanglement

Next to superposition, *entanglement* is one of the most popular and fundamental features in quantum mechanics. Both are crucial elements in quantum algorithms. Whereas superposition can be defined on just one particle, entanglement requires two or more qubits. Precisely, a pair or set of qubits are called entangled when each qubit cannot be described independently of the state of the remaining qubits.

For the simplest example of entanglement we discuss the pure two qubit state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \quad (3.6)$$

This state describes the superposition of two cases: after measurement with probability 0.5, both of the qubits are in the state  $|0\rangle$  or with the same probability both of them are in the state  $|1\rangle$ . When measuring only one qubit, the other qubit will *collapse* into the same state. In this sense the two qubits are entangled.

$|\Psi^+\rangle$  in is one of the four *Bell states* [199] displayed in Figure 3.5b. These states are maximally entangled superpositions of  $|0\rangle$  and  $|1\rangle$ . For all four states yields: measuring one qubit determines the state of the second qubit. The measurement outcomes are *correlated*. Even if the two qubits are spatially separated, this correlation remains, which is discussed in the *EPR paradox* by Albert Einstein, Boris Podolsky, and Nathan Rosen [200].

In general, we name a quantum state of two or more systems *entangled*, if it is not *separable*. Further, we name a pure state of  $n$  subsystems separable if it can be written in the form

$$|\psi\rangle = |\psi_1\rangle \otimes \cdots \otimes |\psi_n\rangle.$$

For mixed states the equivalent form is the convex sum

$$\rho = \sum_i p_i \rho_1^i \otimes \cdots \otimes \rho_n^i.$$

If there exists only a single coefficient with  $p_k \neq 0$ , separable two-particle states are also often called *product states* and can be described in the form  $\rho_{AB} = \rho_A \otimes \rho_B$ .

### State distances

During quantum algorithms, it is often essential to compare two quantum states. We will see a use case in Section 4.2, where the update of a quantum neural network algorithm is based on the difference of two states.

Although many different proposals for measuring how close two quantum states exist, the *fidelity* is one of the most referenced options. We define the fidelity of two density matrices as

$$F(\rho, \sigma) \equiv \left( \text{tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2, \quad (3.7)$$

where  $F(\rho, \rho) = 1$ . The fidelity can be expressed as  $F(|\phi\rangle\langle\phi|, \rho) = \langle\phi|\rho|\phi\rangle$  if one of the states is pure and as  $F(|\psi\rangle, |\phi\rangle) = |\langle\psi|\phi\rangle|^2$  if both of the states are pure. Note that in some literature the fidelity, in contrast to Equation (3.7), is defined differently, namely as  $\bar{F}(\rho, \sigma) \equiv \text{tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}}$ .

If at least one of the states is pure, the fidelity as a measure of closeness has some advantages: It can be measured using the SWAP test, which we will explain later in this work, see Section 4.6. Another essential point is that the fidelity measure has an operational meaning: it defines the probability  $\langle\phi|\rho|\phi\rangle$  (or  $|\langle\psi|\phi\rangle|^2$ ) that  $\rho$  (or  $|\phi\rangle$ ) passes the test being the same as  $|\psi\rangle$  in measurements.

However, when comparing two mixed states, the fidelity still has an interpretation: It is the largest fidelity between any two purifications of given states. Armin Uhlmann describes this fact in [201]: if we assume  $\rho$  and  $\sigma$  are two states of a quantum system  $A$ , and  $B$  is a second system with dimension greater than or equal to the dimension of  $A$ , then the fidelity can be described as

$$F(\rho, \sigma) = \max |\langle\psi_0|\phi_0\rangle|$$

where the maximisation runs over all  $|\psi_0\rangle$  and  $|\phi_0\rangle$  which are purification of  $\rho$  and  $\sigma$  in  $AB$ . Although the authors of [202] show an operational interpretation of this theorem, the drawback of the fidelity becomes clear at this point: due to the maximisation process, excessive computational complexity is required to evaluate the fidelity of two mixed states on a quantum computer.

A good alternative for comparing two mixed states is the *Hilbert-Schmidt* distance [203]

$$d_{\text{HS}}(\rho, \sigma) \equiv \text{tr}((\rho - \sigma)^2),$$

because it can be evaluated on a quantum computer [204, 205]. In this work, we use the Hilbert-Schmidt distance as a training loss with mixed states for graph-structured quantum data, see Section 6.4. It is important to note that this distance reaches its minimum, i.e.  $d_{\text{HS}}(\rho, \rho) = 0$ , if the two compared states coincide. In

### 3. Quantum information

contrast, the fidelity reaches the maximum, the value 1, in this case.

## 3.2 Quantum circuits

To exploit quantum mechanics for quantum computing, so-called *quantum circuits* are needed. These describe the processes of initialising qubit states, applying operations on these qubits and reading out results via measurements. Since we already discussed qubits and measurements, we will start the discussion of quantum circuits by introducing quantum operations, called quantum *gates*. Further, we will explain how these operations and circuits can be depicted. We will conclude this section by discussing exemplary quantum circuits.

### One-qubit gates

Similar to logic gates, which manipulate classical information and are used for classical computing, we define quantum gates. Some of the classical gates have a direct quantum analogue. One example is the *NOT* gate transferring the state  $|0\rangle$  in  $|1\rangle$  and  $|1\rangle$  in  $|0\rangle$ . Written out in Dirac notation this also called Pauli  $X$ -gate is of the form

$$X = |1\rangle\langle 0| + |0\rangle\langle 1|,$$

or in matrix notation with respect to the basis in Equation (3.1)

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

In circuit notation we depict a gate  $X$  as

$$\text{---} \boxed{X} \text{---},$$

where the line denotes the qubit. Applying this gate on the in Equation (3.1) presented basis states gives the expected results, namely

$$\begin{aligned} X|0\rangle &= |1\rangle\langle 0|0\rangle + |0\rangle\langle 1|0\rangle = |1\rangle \\ X|1\rangle &= |1\rangle\langle 0|1\rangle + |0\rangle\langle 1|1\rangle = |0\rangle. \end{aligned}$$

Thus the gate  $X$  interchanges the roles of the two basis states in an arbitrary pure one-qubit state, i.e.

$$X|\psi\rangle = X(\alpha|0\rangle + \beta|1\rangle) = \beta|0\rangle + \alpha|1\rangle = |\psi'\rangle.$$

With the  $X$ -gate as an example, we have shown how one-qubit gates act on pure states. From here, important questions arise: which gates are allowed in quantum computing and how are they constraint? As explained in Section 3.1, every normalised element in  $\mathbb{C}^2$  can describe a pure one-qubit quantum state,

namely  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$  with  $|\alpha|^2 + |\beta|^2 = 1$ . For the state  $U|\psi\rangle = |\psi'\rangle$  resulting from applying the gate  $U$  the same constraint has to be fulfilled, specifically  $\langle\psi|\psi\rangle = 1 = \langle\psi'|\psi'\rangle$ . It follows that the gate has to be *unitary* defined through  $U^\dagger U = 1$ , where  $U^\dagger$  denotes the adjoint of the matrix  $U$ .

Although every unitary matrix is a legitimate quantum gate, some gates are more frequently used in algorithms than others. In Figure 3.2 we list some famous one-qubit gates. The *Hadamard gate*  $H$  provides a possibility to build superposition of states. It is also often applied in order to measure in the basis  $\{|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)\}_\pm$ . In general it is possible to measure in the basis  $\{U|0\rangle, U|1\rangle\}$  by applying  $U^\dagger$  before measuring in the basis  $\{|0\rangle, |1\rangle\}$ .

$$\begin{array}{cccccc}
 \boxed{X} & \boxed{Y} & \boxed{Z} & \boxed{H} & \boxed{T} & \boxed{S} \\
 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} & \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} & \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} & \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} & \begin{pmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{4}} \end{pmatrix} & \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}
 \end{array}$$

Figure 3.2.: **A selection of one-qubit gates.** The gates  $X$  (also called *NOT*),  $Y$  and  $Z$  are represented by Pauli matrices  $\sigma_x$ ,  $\sigma_y$  and  $\sigma_z$ . Further famous gates are the Hadamard gate  $H$ , the  $T$ -gate also known as  $\frac{\pi}{8}$ -gate and the phase gate  $S$ .

Note that all elements of  $SU(2)$ , i.e. those unitaries on  $\mathbb{C}^2$  with  $\det(U) = 1$  are given by  $SU(2) = \{\exp(in\sigma)\}$ , where  $n \in \mathbb{R}^3$  and  $\sigma$  is the vector of Pauli matrices  $\sigma = (\sigma_x, \sigma_y, \sigma_z)$ , see Figure 3.2.

### Multi-qubit gates

So far, we only discussed gates acting on one-qubit states, but also operations on multiple qubits simultaneously, as depicted in Figure 3.3a, are regularly used. An typical example for a multi-qubit gate is the *controlled-not gate*, abbreviated by *CNOT*, see Figure 3.3b. This gate affects two qubits, where one is called the controlled qubit, and the other is the target qubit. If the controlled qubit is in the state  $|1\rangle$ , the basic states of the target qubits get interchanged, i.e. we apply the  $X$  gate. Otherwise, the identity operates.




Figure 3.3.: **Multi-qubit gates.** In general multi-qubit gates are depicted as gates covering diverse lines, symbolising the qubits, see (a). The *CNOT* gate has a special notation, depicted in (b).

### 3. Quantum information

Before we describe the circuits comprised of quantum gates, we want to comment on density matrices in this context. Although the qubits are usually initialised in pure states at the beginning, mixed states occur during quantum circuits. For applying a gate on a mixed state  $\rho$  we multiply the gate from the left side and its adjugate from the right side that is  $\rho' = X\rho X^\dagger$ . Due to unitarity the constraint  $\text{tr}(\rho) = 1 = \text{tr}(\rho')$  stays fulfilled.

#### Circuits

Since it is impossible to directly implement an arbitrary number of different gates in a quantum computer, the question naturally arises whether complex operations can be approximated by breaking them down into sequences of elementary gates. Indeed the Solovay-Kitaev theorem [206] states that every unitary  $U$  can be approximated with an accuracy of  $\epsilon$  by a sequence of gates from  $\mathcal{G}$  of length  $\mathcal{O}(\log^c(\frac{1}{\epsilon}))$ , where  $\mathcal{G}$  is finite universal gate set and  $c$  is a constant based on the choice of  $U$ . The small error is unavoidable in praxis.

Finite sets of quantum gates are called *universal* if it is possible to build any arbitrary single-qubit gate with them. One of the most common universal sets contains the gates  $\{CNOT, H, S, T\}$ . The quantum computers used for the numerics in this work fragment the operations into the gates  $\{CNOT, RZ, SX, X\}$ , where  $RZ = e^{-i\frac{\theta}{2}X}$  and  $S = \sqrt{X}$ . Additionally, the identity operation is usually also directly included in these universal gate sets. The quantum operations built from these sets of gates can be represented in a quantum circuit. Every horizontal line depicts a qubit. Following the lines from left to right depicts the evolution in time.

Next to gates also tensor products, tracing out subsystems and measurements can be depicted in quantum circuits. In Figure 3.4, the circuit describes building the tensor product of two one-qubit states  $\rho_A$  and  $\rho_B$ , applying a gate  $U$  and reducing the system to the first qubit. The resulting state, namely  $\text{tr}_B(U(\rho_A \otimes \rho_B)U^\dagger)$ , is measured.




Figure 3.4.: **Exemplary quantum circuit.** This circuit depicts building the tensor product of  $\rho_A$  and  $\rho_B$ , applying a two-qubit gate  $U$ , tracing out the qubit  $B$  and measuring the resulting state.

Further, it follows from the Stinespring dilation theorem [52] that every description of a quantum circuit through tensor products, unitary transformations and reductions to subsystem is equivalent to the description via a *completely positive* (CP) map. We call a bounded linear operator  $P$  on a Hilbert space *positive* if it a bounded operator  $T$  exists so that  $P = T^\dagger T$ . Using this, the definition of a CP map can be expressed in the following way:



**Definition 3.1** (Completely positive map). *Let  $A$  and  $B$  be  $C^*$ -algebras and  $\mathcal{E} : A \rightarrow B$  a linear map.  $\mathcal{E}$  is called positive if it maps positive elements (corresponding to the positive operators) to positive elements. We can define a new map*

$$id \otimes \mathcal{E} : \mathcal{C}^{n \times n} \otimes A \rightarrow \mathcal{C}^{n \times n} \otimes B$$

and call  $\mathcal{E}$  completely positive if  $id \otimes \mathcal{E}$  is positive for all  $n$ .

Moreover we can write every CP map in the form  $\mathcal{E}(\rho) = \sum_{\alpha} A_{\alpha} \rho A_{\alpha}^{\dagger}$ , where the operators  $A_{\alpha}$  are called *Kraus operators* and satisfy  $\sum_{\alpha} A_{\alpha}^{\dagger} A_{\alpha} = \mathbb{1}$ . This is called *Kraus decomposition* [207].

A more concrete example of a quantum circuit is explained in [52]: the circuit in Figure 3.5a is built of a Hadamard and *CNOT*-gate and outputs the various Bell states, see Figure 3.5b.




Figure 3.5.: **Bell state producing circuit.** The circuit depicted in (a) produces the Bell states according to given input states. The relations between input and output is shown in (b).

### Sampling random unitaries and quantum states

We have already argued that all unitaries are legitimate quantum gates. Therefore, these mathematical operations play a central role in quantum information [52]. When it comes to initialising random unitaries or integrating over all possible unitaries, phrases like “ $U$  is uniformly sampled at random with respect to the Haar measure” are often used in the description of quantum algorithms. In the following, we briefly motivate the *Haar measure* [208]. For a detailed discussion of the *unitary group* we point to [209].

A measure formulates the notion of how objects are distributed. The Haar measure is used to make results of measure theory applicable in group theory. Given a matrix-valued function  $f(Y)$  on the unitary group  $\mathcal{U}(d)$  we write the integral  $\mathcal{U}(d)$  of  $d \times d$  matrices of  $f(Y)$  with respect to Haar measure as

$$I = \int dY f(Y).$$

### 3. Quantum information

The Haar measure is left- and right-invariance with respect to shifts via multiplication. Hence

$$\int dY f(YU) = \int d(Y'U^\dagger) f(Y') = \int dY' f(Y),$$

where  $U \in \mathcal{U}(d)$  is a fixed unitary. The right-invariant can be spelled out, respectively.

Since we can sample unitary operations using the Haar measure, also sampling quantum states uniformly at random is possible: Therefore, we simply generate the unitaries and apply them to a fixed basis state, i.e.  $d|\psi\rangle \equiv dW|0\rangle$  [210]. In this work, the Haar measure is used for determining an optimal lower bound on the probability that a quantum neural network gives an incorrect output for random input, see Chapter 5.

## 3.3 Quantum algorithms

The big goal for quantum algorithms is *quantum supremacy*. This aim is reached when a quantum algorithm can solve a specific task that cannot be solved in any feasible time by today's classical computers. On the one hand, this definition is not very strict and has often spark discussions on whether quantum supremacy was reached in a specific case or not [211]. On the other hand, proposals of such quantum algorithms are sometimes the motivation for new classical algorithms outperforming the proposed quantum algorithm [212].

However, work in this direction already has also shown that quantum circuits can prepare probability distributions that are not reachable in classical computing [104, 213–219]. In this section, we expound on the benefits of quantum algorithms compared to classical ones. Therefore, we study a simplified version of Deutsch's algorithm [52]. In the end, we give a short overview of the most famous early quantum algorithms.

### Exponential memory capacity

One advantage of quantum information is the exponential memory capacity: whereas the state space of a classical computer grows with  $2^b$ , where  $b$  is the number of computational bits, a  $q$ -qubit quantum device is described with a, much larger,  $2^b$ -dimensional Hilbert space, as was explained in Section 3.1. When aiming to solve exponential problems, for example like simulating a quantum system [192, 220], classical computers run out of memory capacity very quick, and quantum computers can lead to opportunities.

### Quantum parallelism

Besides, classical computers compute only classical functions. Assume, for example, a function, which get two bits as an input, i.e. the input is  $\{0, 0\}$ ,  $\{0, 1\}$ ,  $\{1, 0\}$

or  $\{1, 1\}$ . If we want the output for all four inputs, we need to run the function four times. Due to the superposition principle, described in Section 3.1, one can argue that a quantum computer taking a  $q$ -qubit quantum state as an input computes the output for  $2^q$  inputs in parallel. This principle is called *quantum parallelism* [56, 221] and can eventually lead to speed-up over classical computers.

To be more precise, we assume in the following that the aim is to compute outputs of a classical one-bit function  $f(x) : \{0, 1\} \rightarrow \{0, 1\}$  for both inputs, 0 and 1 simultaneously. Two qubits are initialised as  $|x\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}$  and  $|y\rangle = |0\rangle$ . Using a two-qubit unitary  $U_f$ , defined through  $U_f |x, y\rangle = |x, y \oplus f(x)\rangle$ , the function  $f(x)$  can be worked out on both inputs via

$$U_f |x, y\rangle = \frac{|0, f(0)\rangle + |1, f(1)\rangle}{\sqrt{2}}$$

Note that  $f(x)$  is a binary function and  $y \oplus f(x)$  denotes the addition of  $y$  and  $f(x)$  modulo 2.

The problem with quantum parallelism is that measurements are required to extract classical data out of the quantum algorithm. The measurement gives only one result,  $f(0)$  when  $|0\rangle$  is measured and  $f(1)$  for  $|1\rangle$ , respectively. The benefit of parallelism is not available. However, we explain in the following how this problem can be avoided and determining  $f(0) \oplus f(1)$  is possible. To this end, we follow [52] to present a simplified version of Deutsch's algorithm [221], depicted as a quantum circuit in Figure 3.6.

### Deutsch's algorithm

Unlike the ansatz above we initialise two qubits in the state  $|\psi_1\rangle = |xy\rangle$ , where

$$\begin{aligned} |x\rangle &= H |0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \\ |y\rangle &= H |1\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}. \end{aligned}$$

Applying the unitary  $U_f$  gives

$$U_f |\psi_1\rangle = |\psi_2\rangle = \begin{cases} \pm \left( \frac{|0\rangle + |1\rangle}{\sqrt{2}} \right) \otimes \left( \frac{|0\rangle - |1\rangle}{\sqrt{2}} \right) & \text{if } f(0) = f(1) \\ \pm \left( \frac{|0\rangle - |1\rangle}{\sqrt{2}} \right) \otimes \left( \frac{|0\rangle - |1\rangle}{\sqrt{2}} \right) & \text{if } f(0) \neq f(1). \end{cases}$$

For the derivation it helps to see first that

$$U_f \left( |x\rangle \otimes \left( \frac{|0\rangle - |1\rangle}{\sqrt{2}} \right) \right) = (-1)^{f(x)} |x\rangle \otimes \left( \frac{|0\rangle - |1\rangle}{\sqrt{2}} \right).$$

### 3. Quantum information

After applying a Hadamard gate on the first qubit of  $|\psi_1\rangle$  we end with the resulting state

$$(H \otimes \mathbb{1}) |\psi_2\rangle = |\psi_3\rangle = \begin{cases} \pm |0\rangle \otimes \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}}\right) & \text{if } f(0) = f(1) \\ \pm |1\rangle \otimes \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}}\right) & \text{if } f(0) \neq f(1). \end{cases}$$

Since

$$|\psi_3\rangle = \begin{cases} f(0) \oplus f(1) = 0 & \text{if } f(0) = f(1) \\ f(0) \oplus f(1) = 1 & \text{if } f(0) \neq f(1), \end{cases}$$

we can rewrite  $|\psi_3\rangle$  as

$$|\psi_3\rangle = \pm |f(0) \oplus f(1)\rangle \otimes \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}}\right).$$




Figure 3.6.: **Deutsch's algorithm.** The quantum two-qubit circuit for Deutsch's algorithm is build of three Hadamard gates and a Unitary operation  $U_f$ .

Hence when measuring the first qubit of  $|\psi_3\rangle$ , we can determine  $f(0) \oplus f(1)$ , which is remarkable since we only used one evaluation of  $f(x)$  during the process. Although a classical computer can work with probabilistic methods evaluating  $f(0)$  with probability 0.5 and  $f(1)$  with the same probability, a quantum algorithm can work with the interference of the values. In that way, the determination of global information, for example, here  $f(0) \oplus f(1)$ , is possible.

### Famous algorithms

To conclude this section, we want to give a quick overview of the most famous quantum algorithms.

We already motivated and explained Deutsch's algorithm [221] above, which is able to compute whether a one-bit function  $f(x)$  is constant, i.e.  $f(0) \oplus f(1) = 0$  or balanced, i.e.  $f(0) \oplus f(1) = 1$ . The generalisation from a one-bit to n-bit function of Deutsch's algorithm is known as *Deutsch-Jozsa algorithm* [56].

An algorithm that solves a problem of more practical use is *Shor's factoring algorithm* [57]: using the quantum version of the Fourier transform [52, 222] this algorithm can find the prime factors of a given number. Remarkable is that whereas the best nowadays known classical algorithm runs in sub-exponential time [223], Shor's factoring algorithm runs in polynomial time.

Next to quantum Fourier transform-based algorithm, quantum search algorithms gained much attention: the task is to find a particular item given a list of length  $N$ . Classically this problem requires approximately  $N$  operations. Although not providing a speed-up as sensational as Shor’s algorithm, it is noteworthy that *Grover’s search algorithm* [58] solves the task using only approximately  $\sqrt{N}$  operations. More detailed reviews of the mentioned algorithms can be found at [224].

## 3.4 Quantum computers

From the view of a theoretician, it is compelling to sink directly into the theory of new algorithms after learning the basics of quantum computing. However, we should bear in mind that all these approaches have the most beneficial effect in physical realisation. Therefore we will give in the following lines a short introduction to the problems but also the successes of the experimental side of quantum computing.

### Qubit implementation

Since the quantum computer’s elementary component is the qubit, implementing the latter is fundamental for building a quantum computer. As explained in Section 3.1 such a two-level system can be realised through a physical object that can exist in a superposition of two states. The problem is that, on the one hand, these physical qubits have to be well isolated to preserve their features but, on the other hand, accessible for the computational tasks and measurements [52]. Achieving a good balance of both of these requirements is the leading implementation problem. Moreover, a general rule is that the *decoherence time* has to be longer than the gate operation time. The term *decoherence* describes here the irreversible interactions of the qubits with the environment, also referred to as *leakage* [225, 226].

The quality of the qubit implementation can be described by the accuracy with which quantum gates can be performed. The error rate per gate for two-qubit gates with the best hardware for controlling trapped ions [227, 228] or superconducting circuits [229–231] is above 0.1%. This induces a limit in the number of quantum gates that can be applied within one quantum circuit before the noise overwhelms the signal. Other limiting parameters are the time per gate execution and the measurement error probability. The time per gate execution strongly depends on the specific quantum computer model, whereas the measurement error probability is often about 1% [53].

### NISQ devises

Despite these challenges, in the last years, large-scale quantum computation became possible. Big companies like Microsoft, Google and IBM competed in building quantum computers. In 2019, IBM launched a quantum processor containing 53

### 3. Quantum information

qubits [211]. In the same year, Google announced their processor *Sycamore*, also including 53 qubits based on a superconducting circuit [55].

Today’s quantum devices are called *noisy intermediate-scale quantum devices* (NISQ devices) [53–55]. These early quantum computers are characterised by their high sensitivity to the environment (*noisy*) and limited number of available qubits and applicable gates (*intermediate-scale*). Due to these limitations, today’s quantum computers do not achieve fault-tolerant quantum computation [232, 233]. However, future generations of these devices may be able to perform tasks which exceed the capabilities of classical supercomputers.

Noteworthy is moreover that IBM offers the free access to some of their quantum devices [59]. Naturally, the availability of NISQ devices leads to new opportunities. One of them is the implementation of quantum learning algorithms.

## 3.5 Quantum neural networks

In the last section, we discussed the exciting and ongoing progress on quantum computers. Beforehand, Chapter 2 gave an introduction to the widely spreading world of neural networks, including various algorithms, optimisation methods and countless applications. Hence it is not surprising that merging both of these topics in the field of *quantum machine learning* (QML) [60, 61, 68, 70, 73, 234–237] carries great promise for the discovery of new opportunities.

As discussed in Chapter 2, the method of artificial neural networks (NNs) is very popular in classical machine learning (ML). Consequently, in the past years, the construction of a quantum analogue and its usage for QML was of great interest. The term *quantum neural network* (QNN) was introduced first by Subhash Kak in 1995 [238] with the motivation of connecting investigations in the field of neuroscience with characteristics of quantum computation. Whereas the proposals in the ninetenths remain to be of a rather theoretical idea [238–240], today QNNs are mostly viewed as a subset of practical quantum circuits containing parametrised gates. Many different proposals on how to construct these have been made [74, 77, 79–95]. Similar to their classical counterparts QNNs are often build of a fundamental building block. These *quantum perceptrons* have been proposed in various ways in the last years [79, 83–85, 89, 95, 241–245]. Note, that some of the named QNNs are designed for pure quantum tasks, whereas others can be used for classical input. In the latter case, the classical data has to be encoded first. In [246] it is shown, for example, how to encode a binary vector of dimension  $d$  using  $\log_2 d$  qubits.

The detailed discussion of all of these proposals would exceed the scope of this section. However, we want to describe two of them here shortly as they will play a role in the rest of the thesis: Chapter 4 considers comprehensively the QNNs proposed in [79], called *dissipative quantum neural networks* (DQNNs). In this ansatz, the QNN is built of multiple layers of qubits. The quantum perceptrons connect qubits of two consecutive layers with unitary operations and feed-forward the input state over the layers through the network. On the contrary, other QNN

proposals are constructed by just one layer of qubits. The perceptrons are then, for example, defined as a sequence of alternating unitary operators [103–105] which act all on the same qubits. Such an algorithm will be discussed in Section 4.7. Note that an overview of the most recent QNN proposals can be found in [247].

### Implementation

QNNs can be implemented on today’s quantum computers as *variational quantum algorithms* (VQA) [62]. Their process is a quantum-classical hybrid: the algorithms themselves are executed on a quantum computer, but the optimisation process is done classically. In particular *parametrised quantum circuits* [237, 248–250] are used for the implementation. These quantum circuits consist of unitary transformations and can be modified using parameters  $\vec{\theta}$ . Further, a quantum algorithm computes the training loss function. During the training the parameters are updated classically such that the training loss is optimised [101, 248, 251–253], since this task can be efficiently fulfilled in this way.

A substantial benefit of VQAs is that they can be successfully executed on NISQ devices. As explained in Section 3.4 these devices are highly impaired through noise entering with each quantum gate [53]. This limits the number of quantum gates within one quantum circuit before the noise outweighs the algorithm’s performance. It follows that only quantum circuits of small depth can be applied. In contrast to famous methods like Shor’s factoring algorithm [57] or Grover’s search algorithm [58] which require too many gates, QNNs can be executed on these early quantum computers. In Section 4.6 and Section 4.7 we will discuss the implementations of two of the QNN proposals mentioned above.

### Challenges

Although we can train QNNs on today’s NISQ devices, the execution remains challenging: both, the limited number of available qubits as well as the high noise levels can lead to problems and limits. Especially when working with quantum circuits of higher depth, the accurate computation of the loss gradients becomes ambitious [53, 254–256].

This is not the only challenge we have to accept: *Barren plateaus*, the phenomenon of exponentially vanishing gradients of the training loss function in the number of qubits [247, 257–260], can lead to significant problems. The update of the parameters  $\vec{\theta}$  is based on the gradient of the training function, similar to the classical case which was explained in Section 2.4. When this gradient approaches zero for all components the change of the parameters will stagnate and so the training process. Meeting this case, it is advisable to change the architecture of the QNN and/or the training function.

#### Opportunities

As well as their classical analogous, QNNs are able to perform a variety of tasks. The algorithms can be used for tasks including classical data, for example image processing [90, 261]. However, with the advent of quantum computers containing more than just a handful qubits the task of coping with large amounts of quantum data becomes crucial. State tomography [96], or quantum process tomography [97], where the number of needed samples increases exponentially with the number of particles [98] is out of question.

Hence, it is exciting that QNNs allow characterisations with fewer samples compared to tomography methods. Specifically, they can be applied for the classification of classical or quantum data [77, 93, 262] or de-noising quantum data [85, 263–266]. Also the learning of graph-structured quantum data [99, 267–271] was studied. This will be discussed in detail in Chapter 6. Moreover the characterisation of quantum devices is possible, often formulated as the assignment of learning an unknown unitary operation [79, 101, 267, 268, 272, 273].

The latter task will be discussed in the next chapter of this work, where the topic of QNNs will be presented in detail using the example of DQNNs. The discussion includes a precise description of the DQNN architecture and algorithm, implementation details and training results.



# 4

---

## Dissipative quantum neural networks

---

The introductions to *machine learning* (ML) and quantum information in Chapter 2 and Chapter 3 prepared the reader for the following chapter, probably the heart of this thesis. The here presented quantum analogue of a classical neuron [79] leads not only to a quantum feed-forward neural network capable of universal quantum computation with remarkable generalisation behaviour and robustness to noisy training data but is also the basis for further ongoing studies. Two of them, the usage of graph-structured data (see Chapter 6) and an ansatz for a quantum generative adversarial network (see Chapter 7), will be presented later in this work.

As discussed in Chapter 1, one can categorise *quantum machine learning* (QML) techniques into classical ML improving quantum tasks, quantum algorithms speeding up classical ML and exploiting quantum computing devices for tasks with quantum data. The *dissipative quantum neural networks* (DQNNs) is integrated into the latter category. It consists of layers of qubits and can be trained with pairs of quantum states. A training data pair consists of an input state and an aimed output according to the training goal.

The DQNN is built of quantum perceptrons. Such a building block connects two consecutive layers of qubits and can be represented as a completely positive (CP) transition map. The word *dissipative* describes the action of these maps: such a map does not only contain tensoring the state of the current layer to the state of the next layer's qubits and applying unitary operations but also tracing out the qubits from the first of the two layers. In this way, the layer-to-layer transition maps feed-forward input states through the DQNN. The obtained output state can then be compared with the aimed output. This comparison is made through the fidelity of two quantum states and allows conclusions about how the perceptron

## 4. Dissipative quantum neural networks

unitaries must be updated to fulfil the training goal better.

Both, the simulation on classical computers as well as the implementation on quantum computers provide fast optimisation of the network. We focus on supervised training in this chapter and challenge the algorithm with the job of learning an unknown unitary operation. In addition to that, quantum neural networks (QNNs) of this kind can also be exploited for unsupervised learning task and act as *quantum autoencoders* to de-noise entangled quantum states [265].

One advantage of DQNNs is that the dissipative structure leads to reduced memory requirements, since the number of required qubits scales with only the width, not the depth of the QNN. This enables us to train deep QNNs, if the quantum device allows qubits to be “reused”. Furthermore, no disturbing effects of the Barren plateau phenomenon, discussed in Section 3.5, have been observed in the studies of [79]. When the quantum neurons are sufficiently local and sparse, the QNNs might totally avoid these feared vanishing gradients of the training loss function [94]. Finally, numerical results have shown that DQNNs reach the fundamental information-theoretic limits on quantum learning stated in the form of the *quantum no free lunch* (QNFL) theorem. This bound will be discussed in Chapter 5.

This chapter will introduce the reader to DQNNs in their entirety. The first four sections follow the work of [79]. We start by explaining the layer-to-layer transition maps and the general network architecture in Section 4.1. We focus on the task of learning an unknown unitary, hence in Section 4.2 we formulate the training data for this exemplary task and a suitable loss for training the DQNN. Further, we define a second loss function for testing the behaviour of the QNN after training. The training algorithm itself is explained in Section 4.3. All derivations of the rules used for updating the quantum perceptrons are derived at this point. Section 4.4 describes a vivid proof of the above-mentioned universality. The classical simulation of the training leads to promising results, which are presented in Section 4.5. Additionally, the implementation on a NISQ quantum device and training results are discussed in Section 4.6, following [101]. We end the chapter with a comparison of the DQNN algorithm with another QNN architecture in Section 4.7.

### 4.1 Network architecture

In analogy to a classical neural network (NN), see Section 2.2, the DQNN is build of *quantum perceptrons* acting on qubits arranged in layers. We start with a description of the perceptrons which act as building blocks of the QNN and explain how the QNN is built using these afterwards.

#### Quantum perceptron

The perceptrons are engineered as arbitrary unitary operators. Such a perceptron unitary acts on  $m + n$  qubits and depends on  $(2^{m+n})^2 - 1$  parameters. We define




Figure 4.1.: **Quantum perceptron.** The implementation of the perceptron, here acting on 4 qubits with  $m = 3$  and  $n = 1$ , can be depicted in network notation (a). The implementation as quantum circuit includes initialisation of  $m$  qubits in the input state and  $n$  in the zero state, applying a unitary operation  $U$  and tracing out  $m$  qubits (b).

$m$  of the qubits as input qubits and  $n$  as output qubits. It is further needed for the algorithm that the input qubits are initialised in a state  $\rho^{\text{in}}$  and the output qubits in a product state  $|0\dots 0\rangle$ , respectively. After applying the perceptron unitary, the  $m$  input qubits are traced out and we are left with the  $n$ -qubit state  $\rho^{\text{out}}$ . This output state has in general not the same dimension as the input state. Figure 4.1 makes clear how an exemplary perceptron can be depicted in network or quantum circuit notation. One single perceptron can be seen as a small DQNN consisting only of two layers of qubits and one unitary operation.

Before describing how to build a large DQNN constructed of many perceptrons, we want to make two remarks: for simplicity, we set  $n = 1$  here, i.e. the perceptrons are  $m + 1$ -qubit unitaries. As we will see in Section 4.4 this will not threaten the universality of this model. Further, although we suppose working with 2-level qubits, it can be handy to notice that a so defined perceptron can be easily generalised for qudits.

### Quantum neural network

The DQNN is a quantum circuit built out of  $L + 2$  layers of qubits. We name the first layer *input* layer, the last layer *output* layer and count the  $L$  layers in between, namely the *hidden* layers, with the variable  $l$ , see Figure 4.2. In each layer, the perceptron unitaries are applied layerwise and from top to bottom. In that way two subsequent layers are fully connected through perceptrons. Remember that to each perceptron an arbitrary unitary is assigned, and therefore the perceptrons do not generally commute. The order of application of the perceptron unitaries is indicated with over- and under-crossings in Figure 4.2.

In the introduction to this chapter we already mentioned that the whole network can be expressed as a composition of layer-to-layer transition maps. This notation will be highly important through the whole chapter. According to the definition of

#### 4. Dissipative quantum neural networks




Figure 4.2.: **Network representation of a DQNN.** This DQNN consists of  $L+2$  layers of qubits, and  $L + 1$  layers of quantum perceptrons represented by unitary operations  $U_j^l$ .

a single perceptron, we can phrase the output state of the DQNN as

$$\rho^{\text{out}} = \mathcal{E}(\rho^{\text{in}}) = \mathcal{E}^{L+1}(\mathcal{E}^L(\dots \mathcal{E}^2(\mathcal{E}^1(\rho^{\text{in}})) \dots)), \quad (4.1)$$

using the CP maps  $\mathcal{E}^l$  defined via

$$\mathcal{E}^l(X^{l-1}) \equiv \text{tr}_{l-1} \left( \prod_{j=m_l}^1 U_j^l (X^{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|) \prod_{j=1}^{m_l} U_j^{l\dagger} \right), \quad (4.2)$$

where  $U_j^l$  is assigned to the  $j$ th perceptron acting on the qubit layers  $l - 1$  and  $l$ , and  $m_l$  is the total number of perceptrons acting on layers  $l - 1$  and  $l$ . See Definition 3.1 for an introduction to CP maps. From this notation, it becomes clear that the information propagates from the input to the output layer and a quantum feed-forward NN is implemented. This fact is the groundwork for the back-propagation algorithm, which we will discuss in Section 4.3.

We want to close the discussion of the network architecture with a short observation: it is possible to see the quantum circuit of the network as a single unitary  $\mathcal{U} = U^{\text{out}} U^L U^{L-1} \dots U^1$ , where  $U^l = U_{m_l}^l \dots U_1^l$  are the layer unitaries, comprised of a product of quantum perceptrons acting on the qubits in layers  $l - 1$  and  $l$ , see Figure 4.2. However, note that every  $U_k^l$  has to be extended by identities for the remaining qubits to get the correct dimensions when implementing the architecture. We leave these off for a cleaner notation here. Taking this into account the formula

for the output state can be expressed as

$$\mathcal{E}(\rho^{\text{in}}) \equiv \text{tr}_{\text{in, hid}}(\mathcal{U}(\rho^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|)\mathcal{U}^\dagger).$$

## 4.2 Loss functions

After describing the structure of the DQNN and before explaining the algorithm, we need to formulate a training goal. In this section, we do so by formulating two loss functions: one is aimed to be maximised during training, and a second one is used to check the trained DQNN afterwards.

### Training and validation data

We start by specifying a learning task. Therefore, imagine the following setting: we have access to a not-trusted or uncharacterised quantum device or a similar source acting as a unitary  $Y$  on  $m$ -qubit input states. This device can be repeatedly initialised and applied to arbitrary states. In that way it is possible to prepare training data pairs, also called supervised data, structured as  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}$  with  $|\phi_x^{\text{SV}}\rangle = Y|\phi_x^{\text{in}}\rangle$  and  $x = 1, 2, \dots, N$ . Notice that the input state generally does not have to be pure, but we restrict our discussion to this case for simplicity here. In the following we use the notation  $\rho_x^{\text{in}} \equiv |\phi_x^{\text{in}}\rangle \langle \phi_x^{\text{in}}|$  for the input. Notice further that after generating  $N$  training data pairs, we will only use  $S$  of them for training.  $N - S$  pairs will not be used for the training algorithm, but for testing the generalisation behaviour of the trained network as explained later.

At this point we have to mention a crucial point: in our scenario and using the in the afterwards discussed algorithm, it is essential that we can request multiple copies of each of our  $S$  training pairs to overcome quantum projection noise in evaluating the derivative of the loss function. Since *cloning* an unknown quantum state is not possible [52], we can not simply copy the training data pairs. Instead we assume we have a device, where it is possible to press the button  $x$  and we get the training data pair  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}$ . It is worth mentioning at this point that this procedure in combination with QML is still more efficient than characterising a  $n$ -qubit device via tomography [96, 97], where the number of needed samples scales exponentially with the number of qubits  $n$ .

The more neurons we use to build the QNN, the more copies per training round are needed. In detail we need  $n^{\text{proj}} \times n^{\text{par}}$  copies, where  $n^{\text{par}}$  is the total number of parameters in the network and  $n^{\text{proj}}$  the factor coming from repetition of measurements to reduce projection noise.  $n^{\text{par}}$  is linear in the number of neurons. It is given by

$$n^{\text{DQNN}} = \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \#\text{param}(U_j^l) = \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} (4^{(m_{l-1}+1)} - 1) = \sum_{l=1}^{L+1} m_l \times (4^{(m_{l-1}+1)} - 1),$$

where  $\#\text{param}(U_j^l)$  denotes the number of parameters used to describe  $U_j^l$  and  $m_l$

#### 4. Dissipative quantum neural networks

is the number of perceptrons acting on the qubit layers  $l - 1$  and layer  $l$ . The  $-1$  term appears because the overall phase of the unitaries is unimportant. Moreover, we assume that the used DQNN has number of input and output qubits suiting the training data, i.e.  $m = m_0 = m_{L+1}$ .

#### Training loss

Since we already discussed the training data set and goal, we can now investigate the loss functions. Given an architecture in the form of a DQNN, the aim is now to exploit the training pairs to learn the action of the given quantum device and perfectly reproduce it through the network. This means the loss function should not only be meaningful to the problem, but reach a global extremum if the training goal is reached. During the training it is to avoid that the the gradient of the loss function vanishes. Furthermore, we require that the function be efficiently computed on a quantum computer.

The training algorithm, explained in Section 4.3, naturally includes updating the DQNN. This is the centrepiece of the training and the main step to reach our training goal which can be formulated the following way: we desire the network's output  $\mathcal{E}(\rho_x^{\text{in}})$  to be as close as possible to the correct output  $|\phi_x^{\text{SV}}\rangle$  for a specific input  $|\phi_x^{\text{in}}\rangle$ . We use an essentially unique measure of closeness for (pure) quantum states, to quantify this, and define the *training loss function* as the fidelity  $F$  between the QNN output and the desired output, averaged over the training data. We aim to optimise, more precisely to maximise, the loss function during training. For pure supervised states the loss function takes the form

$$\mathcal{L}_{\text{SV}} = \frac{1}{S} \sum_{x=1}^S F(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|, \rho_x^{\text{out}}) = \frac{1}{S} \sum_{x=1}^S \langle \phi_x^{\text{SV}} | \rho_x^{\text{out}} | \phi_x^{\text{SV}} \rangle$$

and has the domain  $[0, 1]$ . We aim for the value 1 during training, since the fidelity reaches 1, if the two compared states are equal.

Choosing the fidelity as a distance measure of a pure state and a mixed state has a few advantages. Not only is the fidelity a good generalisation of the risk function considered in training classical NN, but it can also be measured and has an operational meaning. See Section 3.1 for a more detailed discussion.

#### Validation loss

Quantum data is generally very limited. Therefore it is of great interest to analyse how well the network performs on unseen data. After the training we can do this with checking how well the network predicts the  $N - S$  unsupervised outcomes. For this task we specify the *validation loss*

$$\mathcal{L}_{\text{USV}} = \frac{1}{N - S} \sum_{x=S+1}^N \langle \phi_x^{\text{USV}} | \rho_x^{\text{out}} | \phi_x^{\text{USV}} \rangle.$$

This loss functions also helps to analyse, if the training the DQNN leads to overfitting, i.e. the model matches the training data but does not perform accurately on unseen data, which is generally to avoid.

### 4.3 Training algorithm

With the loss function and QNN ansatz in hand, we can explain how the training proceeds. In the following, we will lead through the algorithm which can be found as a summarised version in Figure 4.3. The rest of this section will be somewhat technical. Nonetheless, we encourage the reader not to skip these parts, since the proofs allow a deeper understanding of the algorithm and structure of the DQNN. We begin with the derivation of the update matrix (see Proposition 4.1). It follows a derivation of the adjoint channel  $\mathcal{F}_t^l$  of  $\mathcal{E}_t^l$  (see Proposition 4.2) as well as a revealing discussion on the form of the update matrix, which can be expressed in terms of the layer-to-layer channels. We close the section with calculating of the change in the training loss during training (see Proposition 4.3) in a way that is similar to the classical back-propagation explained in Section 2.4.

#### Overview

In the first step, the *initialisation*, all perceptron unitaries  $U_j^l$  the DQNN is build of are initialised by randomly chosen unitaries and all qubit states in all layers are initialised in the zero state. The copies of training data pairs are prepared.

The *feed-forward* part of the algorithm promotes the input of every training pair through the network, which is described with Equation (4.1). For every layer, the state of the previous layer  $\rho_x^{l-1}$  is tensored with the layer  $l$  initialised in the state  $|0\dots 0\rangle$ . The unitaries  $U^l = U_{m_l}^l \dots U_1^l$  acting on layer  $l - 1$  and layer  $l$  are applied. Finally, the layer  $l - 1$  is traced out and the resulting state  $\rho_x^l$  is saved and used for the next feed-forward step.

The next piece of the algorithm, the *back-propagation*, is very similar to the previous part. The only differences are that we propagate the *output* part of the training pair, and that the propagation direction is the opposite. This means we start at the end of the network with the state  $\rho_x^{\text{SV}}$  and use the adjoint channel of  $\mathcal{E}^l$  instead of  $\mathcal{E}^l$  itself for the propagation. The resulting states are saved as  $\sigma_x^{l-1}$ . We will derive the exact form of the adjoint channel  $\mathcal{F}_t^l(X^l)$  in Proposition 4.2.

The information gained in the feed-forward and back-propagation parts of the algorithm, saved as states  $\rho_x^l$  and  $\sigma_x^{l-1}$ , is used for *updating* the network's unitaries in a way that optimises the training loss. Assuming that the unitaries depend on a parameter  $t$  and using gradient descent leads to the update rule

$$U_j^l(t + \epsilon) = e^{i\epsilon K_j^l(t)} U_j^l(t),$$

where  $\epsilon$  is the training step size. It follows that also the layer-to-layer map  $\mathcal{E}^l(t) \equiv \mathcal{E}_t^l$  and its adjoint  $\mathcal{F}^l(t) \equiv \mathcal{F}_t^l$  depend on the parameter  $t$ . We will

#### 4. Dissipative quantum neural networks

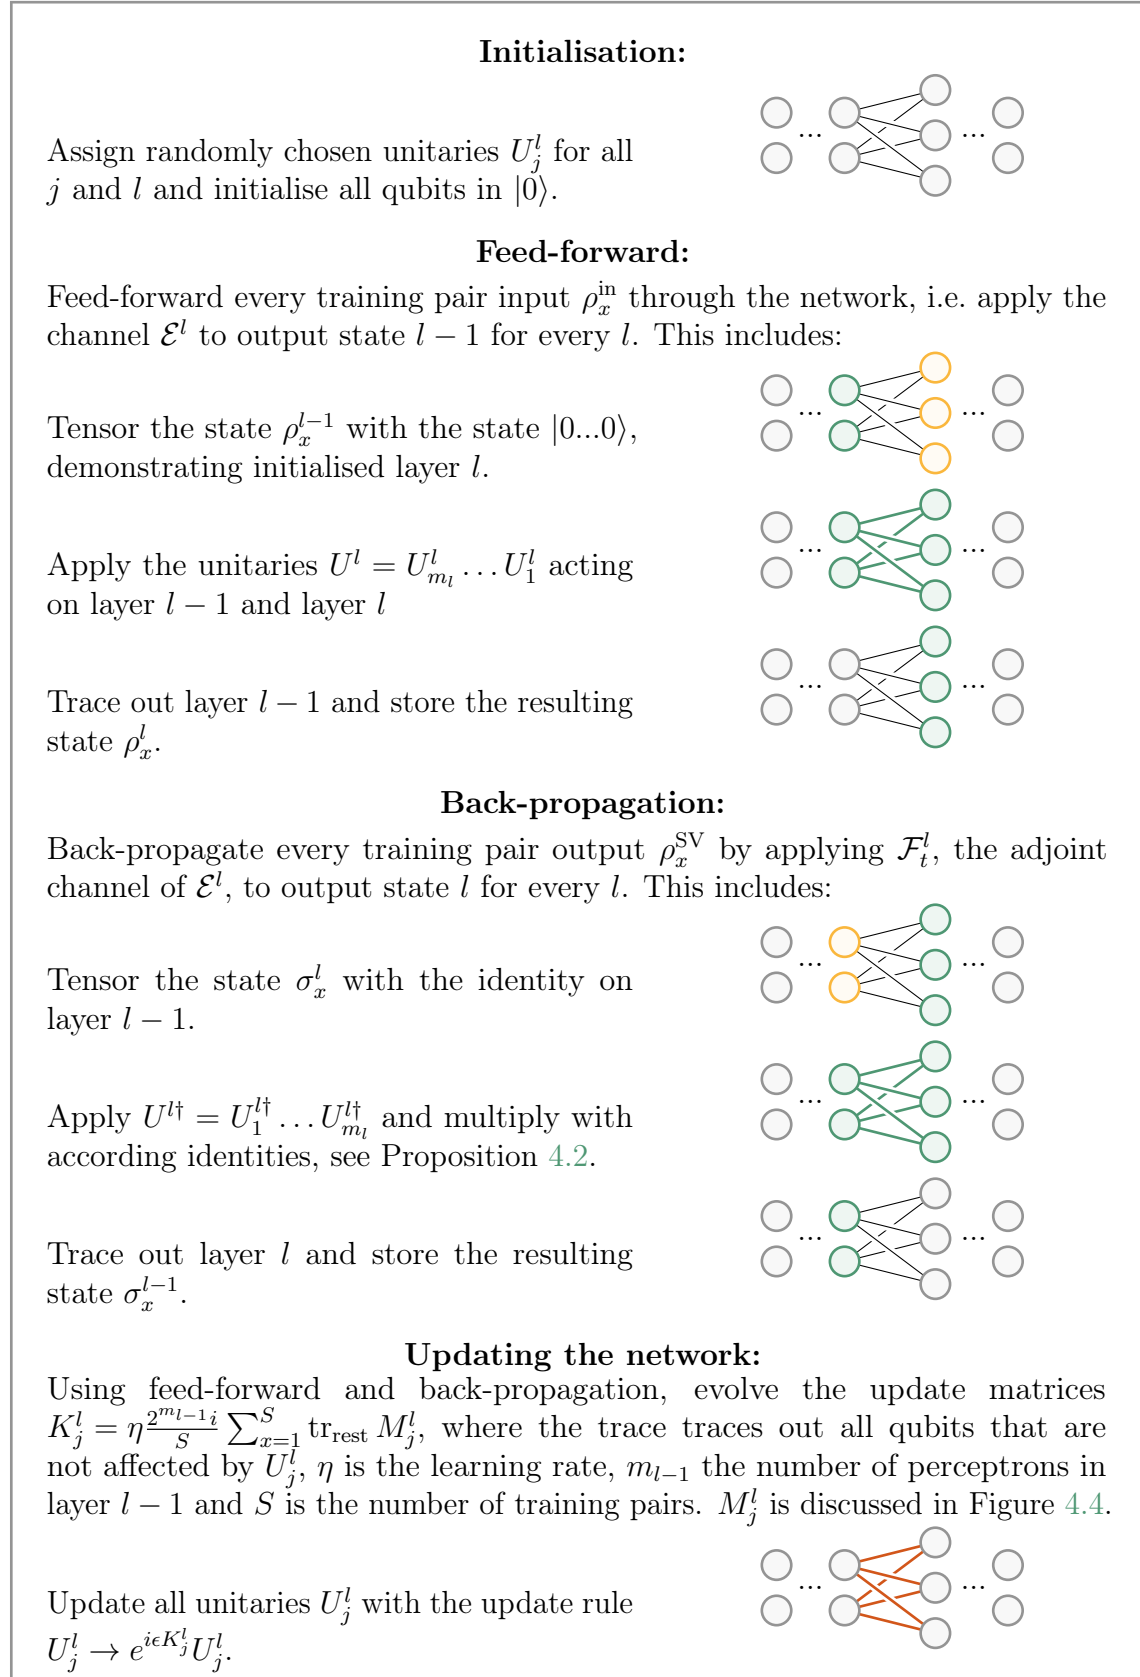


Figure 4.3.: **DQNN training algorithm.** The steps feed-forward, back-propagation and updating are repeated until the training loss reaches its maximum.



determine the update matrix in Proposition 4.1. Note that the steps feed-forward, back-propagation and updating are repeated until the training loss reaches its maximum. We will discuss training results in Section 4.5 (classical simulation) and Section 4.6 (NISQ device implementation).

### Derivation of the update matrix

The above mentioned update rule is based on the update matrices. In the following it will be shown of which form these matrices have to be to improve the perceptrons' performance on the training data.

**Proposition 4.1.** *The update matrix for a QNN trained with pure states  $|\phi_x^{SV}\rangle$  has to be of the form*

$$K_j^l(t) = \frac{\eta 2^{m_l-1} i}{S} \sum_x \text{tr}_{rest} (M_j^l(x, t)),$$

where

$$M_j^l(x, t) = \left[ U_j^l(t) \dots U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger}(t) \dots U_j^{l\dagger}(t), \right. \\ \left. U_{j+1}^l \dagger(t) \dots U_{m_{L+1}}^{L+1} \dagger(t) (\mathbb{1}_{\text{in, hid}} \otimes |\phi_x^{SV}\rangle \langle \phi_x^{SV}|) U_{m_{L+1}}^{L+1}(t) \dots U_{j+1}^l(t) \right],$$

$U_j^l$  is assigned to the  $j$ th perceptron acting on the qubit layers  $l-1$  and  $l$  and  $\eta$  is the learning rate.

*Proof.* To show the statement we derive  $d\mathcal{L}_{\text{SV}}(t)/dt$ . Therefore the output state of the update step  $t + \epsilon$  is needed. Note that the unitaries act on the current qubit layers and we do not write out the identity matrices for a cleaner notation, e.g.  $U_1^2$  is the short form of  $U_1^2 \otimes \mathbb{1}_{2,3,\dots,m_2}$ .

$$\begin{aligned} \rho_x^{\text{out}}(t + \epsilon) &= \text{tr}_{\text{in, hid}} \left( e^{i\epsilon K_{m_{L+1}}^{L+1}(t)} U_{m_{L+1}}^{L+1}(t) \dots e^{i\epsilon K_1^1(t)} U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|) \right. \\ &\quad \left. U_1^{1\dagger}(t) e^{-i\epsilon K_1^1(t)} \dots U_{m_{L+1}}^{L+1} \dagger(t) e^{-i\epsilon K_{m_{L+1}}^{L+1}(t)} \right) \\ &= \rho_x^{\text{out}}(t) + i\epsilon \text{tr}_{\text{in, hid}} \left( [K_{m_{L+1}}^{L+1}(t), U_{m_{L+1}}^{L+1}(t) \dots U_1^1(t) \right. \\ &\quad \left. (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|) U_1^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1} \dagger(t)] + \dots \right. \\ &\quad \left. + U_{m_{L+1}}^{L+1}(t) \dots U_2^1(t) [K_1^1(t), U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|) \right. \\ &\quad \left. U_1^{1\dagger}(t)] U_2^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1} \dagger(t) \right) + \mathcal{O}(\epsilon^2). \end{aligned}$$

The derivative of the loss function, up to the first order in  $\epsilon$ , can be written as

$$\frac{d\mathcal{L}_{\text{SV}}(t)}{dt} = \lim_{\epsilon \rightarrow 0} \frac{\mathcal{L}_{\text{SV}}(t + \epsilon) - \mathcal{L}_{\text{SV}}(t)}{\epsilon}$$

#### 4. Dissipative quantum neural networks

$$\begin{aligned}
&= \lim_{\epsilon \rightarrow 0} \frac{\frac{1}{S} \sum_{x=1}^S \langle \phi_x^{SV} | (\rho_x^{out}(t+\epsilon) - \rho_x^{out}(t)) | \phi_x^{SV} \rangle}{\epsilon} \\
&= \frac{1}{S} \sum_{x=1}^S \text{tr} \left( (\mathbb{1}_{\text{in,hid}} \otimes |\phi_x^{SV}\rangle \langle \phi_x^{SV}|) \left( \left[ iK_{m_{L+1}}^{L+1}(t), U_{m_{L+1}}^{L+1}(t) \dots U_1^1(t) \right. \right. \right. \\
&\quad \left. \left. \left. (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) U_1^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) \right] + \dots \right. \right. \\
&\quad \left. \left. + U_{m_{L+1}}^{L+1}(t) \dots U_2^1(t) \left[ iK_1^1(t), U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) U_1^{1\dagger}(t) \right] \right. \right. \\
&\quad \left. \left. U_2^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) \right) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr} \left( \left[ U_{m_{L+1}}^{L+1}(t) \dots (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) \dots U_{m_{L+1}}^{L+1\dagger}(t), \right. \right. \\
&\quad \mathbb{1}_{\text{in,hid}} \otimes |\phi_x^{SV}\rangle \langle \phi_x^{SV}| \left. \left. \right] iK_{m_{L+1}}^{L+1}(t) + \dots \right. \\
&\quad \left. + \left[ U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) U_1^{1\dagger}(t), \right. \right. \\
&\quad \left. \left. U_2^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) (\mathbb{1}_{\text{in,hid}} \otimes |\phi_x^{SV}\rangle \langle \phi_x^{SV}|) U_{m_{L+1}}^{L+1}(t) \dots U_2^1(t) \right] iK_1^1(t) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr} \left( M_{m_{L+1}}^{L+1}(t) K_{m_{L+1}}^{L+1}(t) + \dots + M_1^1(t) K_1^1(t) \right),
\end{aligned}$$

where

$$\begin{aligned}
M_{m_{L+1}}^{L+1}(t) &\equiv \left[ U_{m_{L+1}}^{L+1}(t) \dots (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) \dots U_{m_{L+1}}^{L+1\dagger}(t), \right. \\
&\quad \left. \mathbb{1}_{\text{in,hid}} \otimes |\phi_x^{SV}\rangle \langle \phi_x^{SV}| \right], \\
M_1^1(t) &\equiv \left[ U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) U_1^{1\dagger}(t), \right. \\
&\quad \left. U_2^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) (\mathbb{1}_{\text{in,hid}} \otimes |\phi_x^{SV}\rangle \langle \phi_x^{SV}|) U_{m_{L+1}}^{L+1}(t) \dots U_2^1(t) \right].
\end{aligned}$$

With  $|0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|$  we denote the initialised qubits of the hidden layers and the output layer. We describe the identity on the space of the input and hidden layer's qubits with  $\mathbb{1}_{\text{in,hid}}$ . With the usage of the Pauli matrices  $\sigma \equiv \{\mathbb{1}, \sigma^x, \sigma^y, \sigma^z\}$ , defined in Figure 3.2, the parameter matrices  $K_j^j$  can be parametrised as

$$K_j^l(t) = \sum_{\alpha_1, \alpha_2, \dots, \alpha_{m_{l-1}}, \beta} K_{j, \alpha_1, \dots, \alpha_{m_{l-1}}, \beta}^l(t) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^{\alpha_{m_{l-1}}} \otimes \sigma^\beta).$$

The index  $\alpha_i$  counts the qubits in the previous layer,  $\beta$  describes the current qubit in layer  $l$ . In the next step, we maximise  $\frac{d\mathcal{L}_{SV}}{dt}$  because we aim to maximise the loss function in every step of the training algorithm. Therefore we introduce a Lagrange multiplier  $\lambda \in \mathbb{R}$  to ensure a finite solution. Without this condition the

extrema would be  $\pm\infty$ .

$$\begin{aligned}
 & \max_{K_{j,\alpha_1,\dots,\beta}^l} \left( \frac{d\mathcal{L}_{\text{SV}}(t)}{dt} - \lambda \sum_{\alpha_i,\beta} K_{j,\alpha_1,\dots,\beta}^l(t)^2 \right) \\
 &= \max_{K_{j,\alpha_1,\dots,\beta}^l} \left( \frac{i}{S} \sum_{x=1}^S \text{tr} (M_{m_{L+1}}^{L+1}(t) K_{m_{L+1}}^{L+1}(t) + \dots + M_1^1(t) K_1^1(t)) \right. \\
 &\quad \left. - \lambda \sum_{\alpha_1,\dots,\beta} K_{j,\alpha_1,\dots,\beta}^l(t)^2 \right) \\
 &= \max_{K_{j,\alpha_1,\dots,\beta}^l} \left( \frac{i}{S} \sum_{x=1}^S \text{tr}_{\alpha_1,\dots,\beta} \left( \text{tr}_{\text{rest}} (M_{m_{L+1}}^{L+1}(t) K_{m_{L+1}}^{L+1}(t) + \dots + M_1^1(t) K_1^1(t)) \right) \right. \\
 &\quad \left. - \lambda \sum_{\alpha_1,\dots,\beta} K_{j,\alpha_1,\dots,\beta}^l(t)^2 \right).
 \end{aligned}$$

The notation  $\text{tr}_{\text{rest}}$  describes tracing out all qubits which are not effected by  $K_{j,\alpha_1,\dots,\beta}^l$ . As a next step we take the derivative of the gained expression with respect to  $K_{j,\alpha_1,\dots,\beta}^l$ :

$$\frac{i}{S} \sum_{x=1}^S \text{tr}_{\alpha_1,\dots,\beta} \left( \text{tr}_{\text{rest}} (M_j^l(t)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \right) - 2\lambda K_{j,\alpha_1,\dots,\beta}^l(t) = 0,$$

This is equivalent to

$$K_{j,\alpha_1,\dots,\beta}^l(t) = \frac{i}{2S\lambda} \sum_{x=1}^S \text{tr}_{\alpha_1,\dots,\beta} \left( \text{tr}_{\text{rest}} (M_j^l(t)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \right).$$

We finally can express the parameter matrices as

$$\begin{aligned}
 K_j^l(t) &= \sum_{\alpha_1,\dots,\beta} K_{j,\alpha_1,\dots,\beta}^l(t) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \\
 &= \frac{i}{2S\lambda} \sum_{\alpha_1,\dots,\beta} \sum_{x=1}^S \text{tr}_{\alpha_1,\dots,\beta} \left( \text{tr}_{\text{rest}} (M_j^l(t)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \right) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \\
 &= \frac{\eta 2^{m_l-1} i}{S} \sum_{x=1}^S \text{tr}_{\text{rest}} (M_j^l(t)),
 \end{aligned}$$

where  $\eta = 1/\lambda$  is the learning rate and

$$\begin{aligned}
 M_j^l(t) &= \left[ U_j^l(t) U_{j-1}^l(t) \dots U_1^l(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid,out}} \langle 0\dots 0|) U_1^{l\dagger}(t) \dots U_{j-1}^{l\dagger}(t) U_j^{l\dagger}(t), \right. \\
 &\quad \left. U_{j+1}^{l\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) (\mathbb{1}_{\text{in,hid}} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) U_{m_{L+1}}^{L+1}(t) \dots U_{j+1}^l(t) \right].
 \end{aligned}$$

□

#### 4. Dissipative quantum neural networks

##### Derivation of the adjoint channel

Knowing how the update is done, we can describe the entire training algorithm. Nevertheless, we aim to understand the structure of the update matrices more deeply. To write the commutator  $M_j^l(t)$  in terms of layer-to-layer channels, we need the adjoint channel of  $\mathcal{E}$  which is derived in the following.

**Proposition 4.2.** *The adjoint channel  $\mathcal{F}_t(X)$  for the CP map  $\mathcal{E}_t^l(X^{l-1}) = \text{tr}_{l-1}(U^l(t)(X^{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|)U^{l\dagger}(t))$  is of the form*

$$\mathcal{F}_t^l(X^l) = \text{tr}_l((\mathbb{1}_{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|)U^{l\dagger}(t)(\mathbb{1}_{l-1} \otimes X^l)U^l(t)).$$

*Proof.* In order to show which form the adjoint channel  $\mathcal{F}_t^l$  has, we express  $\mathcal{E}_t^l$  in its Kraus representation [207], which was introduced in Section 3.2. For any operator  $X^{l-1}$  on the  $(l-1)$ th layer we can phrase

$$\mathcal{E}_t^l(X^{l-1}) = \sum_{\alpha} A_{\alpha}^l(t)X^{l-1}A_{\alpha}^{l\dagger}(t),$$

where the Kraus operators  $A_{\alpha}$  are maps from the  $(l-1)$ th layer of  $m_{l-1}$  qubits to the  $l$ th layer of  $m_l$  qubits. Naturally the adjoint channel  $\mathcal{F}_t^l$  can be written as

$$\mathcal{F}_t^l(X^l) = \sum_{\alpha} A_{\alpha}^{l\dagger}(t)X^l A_{\alpha}^l(t), \quad (4.3)$$

for any operator  $X^l$  on the  $l$ th layer. Our aim is to express the Kraus operators  $A_{\alpha}^l$ . We choose  $\{|\alpha\rangle\}_{\alpha}$  to be an orthonormal basis in the  $(l-1)$ th layer. Further we assume  $|b\rangle, |c\rangle$  are any vectors in the  $(l-1)$ th layer and  $|d\rangle, |e\rangle$  any vectors in the  $l$ th layer. Using Equation (4.2) and  $U^l(t) = U_{m_l}^l(t) \dots U_1^l(t)$  we get

$$\begin{aligned} \langle d | \mathcal{E}_t^l(|b\rangle \langle c|) | e \rangle &= \langle d | \text{tr}_{l-1}(U^l(t)(|b\rangle \langle c| \otimes |0\dots 0\rangle_l \langle 0\dots 0|)U^{l\dagger}(t)) | e \rangle \\ &= \sum_{\alpha} \langle \alpha, d | U^l(t)(|b\rangle \langle c| \otimes |0\dots 0\rangle_l \langle 0\dots 0|)U^{l\dagger}(t) | \alpha, e \rangle \\ &= \sum_{\alpha} \langle \alpha, d | U^l(t)|b, 0\dots 0\rangle \langle n, 0\dots 0 | U^{l\dagger}(t) | \alpha, e \rangle. \end{aligned}$$

Defining  $A_{\alpha}^l(t)$  with  $\langle d | A_{\alpha}^l(t) | b \rangle = \langle \alpha, d | U^l(t) | b, 0\dots 0 \rangle$  and using Eq. (4.3) we reach the expression

$$\begin{aligned} \langle b | \mathcal{F}_t^l(|d\rangle \langle e|) | c \rangle &= \sum_{\alpha} \langle b | A_{\alpha}^{l\dagger}(t) | d \rangle \langle e | A_{\alpha}^l(t) | c \rangle \\ &= \sum_{\alpha} \langle b, 0\dots 0 | U^{l\dagger}(t) | \alpha, i \rangle \langle \alpha, e | U^l(t) | c, 0\dots 0 \rangle \\ &= \langle b, 0\dots 0 | U^{l\dagger}(t)(\mathbb{1}_{l-1} \otimes |d\rangle \langle e|)U^l(t) | c, 0\dots 0 \rangle \\ &= \langle b | \text{tr}_l(\mathbb{1}_{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0| U^{l\dagger}(t)(\mathbb{1}_{l-1} \otimes |d\rangle \langle e|)U^l(t)) | c \rangle. \end{aligned}$$

Extracting the action of  $\mathcal{F}^l(t)$  on a general operator  $X^l$  leads to

$$\mathcal{F}_t^l(X^l) = \text{tr}_l (\mathbb{1}_{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|_l U^{l\dagger}(t) (\mathbb{1}_{l-1} \otimes X^l) U^l(t)).$$


□

### Beneficial form of the update matrix

We successfully derived the adjoint channel of  $\mathcal{E}$ . At this point, we can discuss the structure of

$$M_j^l(x, t) = \left[ U_j^l(t) \dots U_1^l(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{l\dagger}(t) \dots U_j^{l\dagger}(t), \right. \\ \left. U_{j+1}^{l\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) (\mathbb{1}_{\text{in, hid}} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) U_{m_{L+1}}^{L+1}(t) \dots U_{j+1}^l(t) \right],$$

see Proposition 4.1, in more detail. At a second glance, the commutator expression presents an exciting structure: the first part of the commutator is the input state propagated through the network until we reach the unitary  $U_j^l$  we wish to update. The second part is obtained by back-propagation of the matching desired supervised output state with stopping right before  $U_j^l$ . This observation is visualised in Figure 4.4.



(a) First element of the commutator.

(b) Second element of the commutator.

Figure 4.4.: **The commutator**  $M_j^l$ . The actions of the two parts of  $M_j^l$  become clear in network notation.

Using the the derived channels we can rewrite  $M_j^l(x, t)$  as

$$M_j^l(x, t) = \left[ U_j^l(t) \dots U_1^l(t) (\rho_x^{l-1}(t) \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{l\dagger}(t) \dots U_j^{l\dagger}(t), \right. \\ \left. U_{j+1}^{l\dagger}(t) \dots U_{m_l}^{l\dagger}(t) (\mathbb{1}_{\text{in, hid}} \otimes \sigma_x^l(t)) U_{m_l}^l(t) \dots U_{j+1}^l(t) \right],$$

#### 4. Dissipative quantum neural networks

where the density matrix of the  $l$  layer concerning the  $x$ th training data can be expressed via  $\rho_x^{l-1}(t) = \mathcal{E}_t^{l-1}(\dots \mathcal{E}_t^2(\mathcal{E}_t^1(\rho_x^{\text{in}})) \dots)$ . In an analogous way we write  $\sigma_x^l(t) = \mathcal{F}_t^{l+1}(\dots \mathcal{F}_t^L(\mathcal{F}_t^{L+1}(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|)) \dots)$  using back-propagation. Note here again that both, the feed-forward channels  $\mathcal{E}$  as well as the corresponding adjoint channels, depend on the unitary parameter  $t$  and change during the training.

At this point, it becomes clear why the layer structure of the network is beneficial: to update a perceptron unitary  $U_j^l$ , we need to evaluate  $K_j^l$ . Therefore only the output state of the previous layer,  $\rho^{l-1}$ , obtained by feed-forward propagation through the network, and the state of the following layer  $\sigma^l$ , obtained by back-propagation of the desired output up to the current layer is needed. It follows that the parameter matrices can be received with only accessing two qubit layers at any time, and there is no need to access the whole network for updating a single perceptron. This fact allows us to train deep DQNNs.

### Change of the training loss

Connecting to Section 2.4 where the classical back-propagation algorithm is described, we want to close this technical section with the derivation of a handy formula for the change in the training loss function using the layer-to-layer channels.

**Proposition 4.3.** *The change in the loss function can be written as*

$$\frac{d\mathcal{L}_{\text{SV}}(t)}{dt} = \frac{i}{S} \sum_{x=1}^N \sum_{l=1}^{L+1} \text{tr} \left( \sigma_x^l(t) \mathcal{D}_t^l(\rho_x^{l-1}(t)) \right),$$

where  $\mathcal{D}_t^l = \partial \mathcal{E}_t^l / \partial t$  is the derivative of the channel  $\mathcal{E}_t^l$ .

*Proof.* We already evaluated  $\frac{d\mathcal{L}_{\text{SV}}(t)}{dt}$  in the proof of Proposition 4.1. As a primary step we translate the gained expressions to the channel formalism and get

$$\begin{aligned} \frac{d\mathcal{L}_{\text{SV}}(t)}{dt} &= \frac{i}{S} \sum_{x=1}^S \text{tr} \left( \mathbb{1}_{\text{in, hid}} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}| \right. \\ &\quad \left( \left[ K_{m_{L+1}}^{L+1}(t), U_{m_{L+1}}^{L+1}(t) \dots U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|) U_1^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) \right] \right. \\ &\quad \left. + \dots + U_{m_{L+1}}^{L+1}(t) \dots U_2^1(t) \left[ K_1^1(t), U_1^1(t) (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|) U_1^{1\dagger}(t) \right] \right. \\ &\quad \left. U_2^{1\dagger}(t) \dots U_{m_{L+1}}^{L+1\dagger}(t) \right) \\ &= \frac{i}{S} \sum_{x=1}^S \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( \left( \mathbb{1}_{\text{in, hid}} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}| \right) U_{m_{L+1}}^{L+1} \dots U_{j+1}^l \left[ K_j^l, \right. \right. \\ &\quad \left. \left. U_j^l \dots U_1^1 (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{\text{hid, out}} \langle 0\dots 0|) U_1^{1\dagger} \dots U_j^{l\dagger} \right] U_{j+1}^{l\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) \end{aligned}$$

### 4.3. Training algorithm

$$\begin{aligned}
&= \frac{i}{S} \sum_{x=1}^N \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( U_1^{l+1\dagger}(t) \dots U_{m_{L+1}}^{l+1\dagger}(t) (\mathbb{1}_L \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) U_{m_{L+1}}^{L+1}(t) \dots U_1^{l+1}(t) \right. \\
&\quad \left. U_{m_j}^l(t) \dots U_{j+1}^l(t) \left[ K_j^l(t), U_j^l(t) \dots U_1^l(t) (\rho_x^{l-1} \otimes |0\dots 0\rangle_{l+1} \langle 0\dots 0|) U_1^{l\dagger}(t) \dots U_j^{l\dagger}(t) \right] \right. \\
&\quad \left. U_{j+1}^{l\dagger}(t) \dots U_{m_j}^{l\dagger}(t) \right).
\end{aligned}$$

For a shorter notation we define

$$\begin{aligned}
A &= U_x^{l+1\dagger} \dots U_{L+1}^{l+1\dagger} (\mathbb{1}_{l,\dots,L} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) U_x^{L+1} \dots U_x^{l+1} \\
B &= U_{m_l}^l \dots U_{j+1}^l \left[ K_j^l, U_j^l \dots U_1^l (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{1,\dots,l} \langle 0\dots 0|) U_1^{1\dagger} \dots U_j^{l\dagger} \right] U_{j+1}^{l\dagger} \dots U_{m_l}^{l\dagger}
\end{aligned}$$

and rewrite the change of the loss function as

$$\begin{aligned}
\frac{d\mathcal{L}_{\text{SV}}}{dt} &= \frac{i}{S} \sum_{x=1}^S \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( (\mathbb{1}_{0,\dots,l-1} \otimes A) (B \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( A (\text{tr}_{0,\dots,l-1}(B) \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( A (\text{tr}_{0,\dots,l-1}(B) \otimes \mathbb{1}_{l+1,\dots,L+1}) (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) A (\text{tr}_{0,\dots,l-1}(B) \otimes \mathbb{1}_{l+1,\dots,L+1}) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \sum_{l=1}^{L+1} \sum_{j=1}^{m_l} \text{tr} \left( \text{tr}_{l+1,\dots,L+1} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) A \right) \text{tr}_{0\dots l-1}(B) \right).
\end{aligned} \tag{4.4}$$

After propagating the state  $\rho_x^{\text{in}}$  through the first  $l-1$  qubit layers of the network we get the expression

$$\begin{aligned}
\rho_x^{l-1} &= \mathcal{E}^{l-1}(\dots \mathcal{E}^1(\rho_x^{\text{in}}) \dots) \\
&= \text{tr}_{l-2} \left( (U^{l-1} (\mathcal{E}^{l-2}(\dots \mathcal{E}^{l-3}(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) \dots) \otimes \mathbb{1}_{l-1}) U^{l-1\dagger}) \right) \\
&= \dots \\
&= \text{tr}_{1,\dots,l-2} \left( U^{l-1} \dots U^1 (\mathcal{E}^1(\rho_x^{\text{in}}) \otimes |0\dots 0\rangle_{2,\dots,l-1} \langle 0\dots 0|) U^{2\dagger} \dots U^{l-1\dagger} \right) \\
&= \text{tr}_{1,\dots,l-2} \left( U^{l-1} \dots U^2 \left( \text{tr}_0 \left( U^1 (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_1 \langle 0\dots 0|) U^{1\dagger} \right) \otimes |0\dots 0\rangle_{2,\dots,l-1} \langle 0\dots 0| \right) \right. \\
&\quad \left. U^{2\dagger} \dots U^{l-1\dagger} \right) \\
&= \text{tr}_{0,\dots,l-2} \left( U^{l-1} \dots U^1 (\rho_x^{\text{in}} \otimes |0\dots 0\rangle_{1,\dots,l-1} \langle 0\dots 0|) U^{1\dagger} \dots U^{l-1\dagger} \right).
\end{aligned}$$

#### 4. Dissipative quantum neural networks

This can be used to express  $\text{tr}_{0,\dots,l-1}(B)$  via

$$\begin{aligned}
\text{tr}_{0,\dots,l-1}(B) &= \text{tr}_{0,\dots,l-1} \left( U_{m_l}^l \dots U_{j+1}^l \left[ K_j^l, U_j^l \dots U_1^l \left( \rho_x^{\text{in}} \otimes |0\dots 0\rangle_{1,\dots,l} \langle 0\dots 0| \right) U_1^{1\dagger} \dots U_j^{l\dagger} \right] \right. \\
&\quad \left. U_{j+1}^{l\dagger} \dots U_{m_l}^{l\dagger} \right) \\
&= \text{tr}_{l-1} \left( U_{m_l}^l \dots U_{j+1}^l \left[ K_j^l, U_j^l \dots U_1^l \left( \text{tr}_{0,\dots,l-2} \left( U_x^{l-1} \dots U^1 \right. \right. \right. \right. \\
&\quad \left. \left. \left. \left( \rho_x^{\text{in}} \otimes |0\dots 0\rangle_{1,\dots,l-1} \langle 0\dots 0| \right) U_1^{1\dagger} \dots U^{l-1\dagger} \right) \right. \right. \\
&\quad \left. \left. \left. \left. \otimes |0\dots 0\rangle_l \langle 0\dots 0| \right) U_1^{l\dagger} \dots U_j^{l\dagger} \right] U_{j+1}^{l\dagger} \dots U_{m_l}^{l\dagger} \right) \\
&= \text{tr}_{l-1} \left( U_{m_l}^l \dots U_{j+1}^l \left[ K_j^l, U_j^l \dots U_1^l \left( \rho_x^{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0| \right) U_1^{l\dagger} \dots U_j^{l\dagger} \right] \right. \\
&\quad \left. U_{j+1}^{l\dagger} \dots U_{m_l}^{l\dagger} \right). \tag{4.5}
\end{aligned}$$

In the proof of Proposition 4.2 it was shown that the structure of the adjoint channel of  $\mathcal{E}^l$  is

$$\mathcal{F}^l(X^l) = \text{tr}_l \left( (\mathbb{1}_{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|) U^{l\dagger} (\mathbb{1}_{l-1} \otimes X^l) U^l \right).$$

With this in hand we can back propagate a supervised state  $|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|$   $L-l$  qubit layers through the network and ending up with the state

$$\begin{aligned}
\sigma_x^l &= \mathcal{F}^{l+1} \left( \dots \mathcal{F}^{L+1} \left( |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}| \right) \dots \right) \\
&= \text{tr}_{l+1} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1} \langle 0\dots 0|) U^{l+1\dagger} \left( \mathbb{1}_l \otimes \mathcal{F}^{l+2} \left( \dots \mathcal{F}^{L+1} \left( |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}| \right) \dots \right) \right) U^{l+1} \right) \\
&= \dots \\
&= \text{tr}_{l+1,\dots,L} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L} \langle 0\dots 0|) \right. \\
&\quad \left. U^{l+1\dagger} \dots U^{L\dagger} \left( \mathbb{1}_{l,\dots,L-1} \otimes \mathcal{F}^{L+1} \left( |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}| \right) \right) U^L \dots U^{l+1} \right) \\
&= \text{tr}_{l+1,\dots,L} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L} \langle 0\dots 0|) \right. \\
&\quad \left. U^{l+1\dagger} \dots U^{L\dagger} \left( \mathbb{1}_{l,\dots,L-1} \otimes \text{tr}_{L+1} \left( (\mathbb{1}_L \otimes |0\dots 0\rangle_{L+1} \langle 0\dots 0|) \right. \right. \right. \\
&\quad \left. \left. \left. U^{L+1\dagger} \left( \mathbb{1}_L \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}| \right) U^{L+1} \right) \right) U^L \dots U^{l+1} \right) \\
&= \text{tr}_{l+1,\dots,L+1} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L} \langle 0\dots 0| \otimes \mathbb{1}_{L+1}) \right. \\
&\quad \left. (U^{l+1\dagger} \dots U^{L\dagger} \otimes \mathbb{1}_{L+1}) (\mathbb{1}_{l,\dots,L} \otimes |0\dots 0\rangle_{L+1} \langle 0\dots 0|) \right. \\
&\quad \left. (\mathbb{1}_{l,\dots,L-1} \otimes U^{L+1\dagger}) (\mathbb{1}_{l,\dots,L} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) (\mathbb{1}_{l,\dots,L-1} \otimes U^{L+1}) (U^L \dots U^{l+1} \otimes \mathbb{1}_{L+1}) \right) \\
&= \text{tr}_{l+1,\dots,L+1} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L} \langle 0\dots 0| \otimes \mathbb{1}_{L+1}) (\mathbb{1}_{l,\dots,L} \otimes |0\dots 0\rangle_{L+1} \langle 0\dots 0|) \right)
\end{aligned}$$



### 4.3. Training algorithm

$$\begin{aligned}
& (U_x^{l+1\dagger} \dots U^{L\dagger} \otimes \mathbb{1}_{L+1}) (\mathbb{1}_{l,\dots,L-1} \otimes U^{L+1\dagger}) (\mathbb{1}_{l,\dots,L} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) (\mathbb{1}_{l,\dots,L-1} \otimes U^{L+1}) \\
& (U^L \dots U^{l+1} \otimes \mathbb{1}_{L+1}) \\
& = \text{tr}_{l+1,\dots,L+1} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) U^{l+1\dagger} \dots U^{L+1\dagger} (\mathbb{1}_{l\dots L} \otimes |\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) \right. \\
& \quad \left. U^{L+1} \dots U^{l+1} \right) \\
& = \text{tr}_{l+1,\dots,L+1} \left( (\mathbb{1}_l \otimes |0\dots 0\rangle_{l+1,\dots,L+1} \langle 0\dots 0|) A \right).
\end{aligned}$$

With this, Equation (4.4) and Equation (4.5) we acquire the expression

$$\begin{aligned}
\frac{d\mathcal{L}_{\text{SV}}(t)}{dt} &= \frac{i}{S} \sum_{x=1}^N \sum_{l=1}^{L+1} \text{tr} \left( \sigma_x^l \sum_{j=1}^{m_j} U_{m_j}^l(t) \dots U_{j+1}^l(t) [K_j^l(t), U_j^l(t) \dots U_1^l(t) \right. \\
& \quad \left. (\rho_x^{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|) U_1^{l\dagger}(t) \dots U_j^{l\dagger}(t) \right] U_{j+1}^{l\dagger}(t) \dots U_{m_j}^{l\dagger}(t) \left. \right).
\end{aligned}$$

It is possible to reformulate the latter equation as

$$\frac{d\mathcal{L}_{\text{SV}}(t)}{dt} = \frac{i}{S} \sum_{x=1}^N \sum_{l=1}^{L+1} \text{tr} \left( \sigma_x^l(t) \mathcal{D}_t^l(\rho_x^{l-1}(t)) \right),$$

where  $\mathcal{F}_t^l$  denotes the adjoint channel of the channel  $\mathcal{E}_t^l$  and the quantum state  $\sigma_x^l(t)$  is defined as  $\sigma_x^l(t) = \mathcal{F}_t^{l+1}(\dots \mathcal{F}_t^{\text{out}}(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|) \dots)$ . Further  $\mathcal{D}_t^l = \partial \mathcal{E}_t^l / \partial t$  is the derivative of the channel  $\mathcal{E}_t^l$ , defined as

$$\begin{aligned}
\mathcal{D}_t^l(X^{l-1}) &= \sum_{j=1}^{m_j} \text{tr}_{l-1} \left( U_{m_j}^l \dots U_{j+1}^l [K_j^l, U_j^l \dots U_1^l (\rho_x^{l-1} \otimes |0\dots 0\rangle_l \langle 0\dots 0|) U_1^{l\dagger} \dots U_j^{l\dagger} \right. \\
& \quad \left. U_{j+1}^{l\dagger} \dots U_{m_j}^{l\dagger} \right).
\end{aligned}$$

□

Due to the length of this section, we take a moment to conclude the results. In the preceding pages, we explained the training algorithm of the DQNN. We not only derived the desired necessary update rule for optimising training loss, defined Section 4.2, but also showed that for the update of the  $l$ th layer of the network, only the quantum states of two neighbouring layers are needed. Finally, we derived an expression for the change of the mentioned loss during training in equivalence to the classical back-propagation. Before we proceed with presenting the DQNN architecture with convincing numerical results, we discuss the previously mentioned universality of the DQNN.

## 4.4 Universality

Classical NNs composed of classical perceptrons can represent any function [274]. In the following, we discuss that the in Section 4.1 presented quantum perceptron holds the same feature for QNNs.

As already motivated in Section 3.2, from the Stinespring dilation theorem [52] directly follows that every description of a quantum circuit through tensor products, unitary transformations and reductions to subsystems is equivalent to the description as a CP map. Figure 4.5 describes this relation using a quantum circuit. The only task is to find the suiting unitary. This is exactly what the algorithm presented in Section 4.3 does. Therefore it is clear that we can implement any CP map with the DQNN algorithm.




Figure 4.5.: **Stinespring theorem in quantum circuits.** A CP map  $\mathcal{E}$  can be represented using tensoring an ancilla state, applying a unitary transformation and tracing out a subsystem.

However, it is more remarkable that a QNN, comprised of the DQNN quantum perceptrons acting on 4-level qudits (equivalent to pairs of 2-level qubits) that commute within each layer, is capable of carrying out universal quantum computation. In the following, we use such a construction to illustrate that it is possible to construct a DQNN simulating an arbitrary quantum circuit consisting out of these perceptrons. Without loss of generality only two qubit gates on neighbouring qubits are used in the circuit.

To construct a QNN that is equivalent to that circuit, we first number the neurons of the QNN, the 4-level qudits, by two indices. Neuron  $(l, j)$  is the  $j$ th neuron in  $l$ th layer. We assume there are  $m_l$  neurons in the  $l$ th layer. Further we suppose that every neuron  $(l, j)$  is connected to neurons  $(l-1, j)$  and  $(l+1, j - (-1)^l \bmod m_l)$  and other connections do not exist. The neurons are marked as rectangles in Figure 4.6a.

We further construct the neurons in a way that every neuron corresponds to two qubits. Each neurons qubits are labelled  $A$  and  $B$  and initialised in the state  $|00\rangle$ .




Figure 4.6.: **DQNN universality proof.** A DQNN constructed out of pairs of perceptrons using unitary operations  $U_j^l$  and *SWAP* gates (a) is equivalent to a quantum circuit of two-qubit gates  $U_j^l$  acting on alternating pairs of qubits (b).

We consider that every operation of the network has the form

$$\rho^l = \text{tr}_{l-1}(U^{l-1}(\rho^{l-1} \otimes [\bigotimes_{j=0}^{m_l-1} |00\rangle_{(l,j)} \langle 00|])U^{l-1\dagger}).$$

$U^l = \prod_{j=m_l-1}^0 U_j^l$  consists of the perceptrons acting on the qubit layers  $l$  and  $l+1$ . Note that, different to the preceding chapters, we let the indices  $l$  and  $j$  start at 0. This results in a much simpler expression of  $U_j^l$  when using the modulo notation. The operation acting on the neuron  $(l, j)$  are defined as

$$U_j^l = V_j^l \text{SWAP}[(l, j, A), (l+1, j - \frac{1 - (-1)^l}{2} \bmod m_l, B)] \\ \text{SWAP}[(l, j, B), (l+1, j + \frac{1 + (-1)^l}{2} \bmod m_l, A)],$$

#### 4. Dissipative quantum neural networks

see Figure 4.6a. Notice that the SWAP operators acting on two qubits: on one qubit that corresponds to the neuron  $(l, j)$  and one of the  $l + 1$ th layer.  $V_j^l$  is a unitary that acts on the qubits of the neuron  $(l, j)$ . All the SWAP operations commute for a fixed  $l$ , because they act on different pairs of qubits.

An in that way constructed DQNN is equivalent to the quantum circuit of two-qubit gates  $V_j^l$  that act on alternating pairs of qubits like depicted in Figure 4.6b. Qubit  $A$  in neuron  $(0, 0)$  on the left hand side diagram corresponds to the qubit labelled by  $a$  on the right hand side. Similarly, the  $(0, 0)$   $B$  qubit corresponds to the qubit labelled by  $b$  and so on.

It is known that two-qubit gates are universal [52]. The SWAP operation is a two-qubit operation, thus the pictured quantum circuit is universal. It is clear that the construction depicted in Figure 4.6 is by far not the most efficient ones but shows the universality of the in a very neat way.

### 4.5 Classical simulation

Different aspects of the DQNN training algorithm have been discussed in this chapter so far. However, no training results were presented heretofore. The algorithm described in Section 4.3 can be indeed fully simulated on a classical computer. Note that since the Hilbert space dimension scales exponentially with the number of qubits, the simulation is even on supercomputers restricted to a few qubits. The simulation is still useful to study the behaviour of the algorithm.

Hence in this section, we describe a classical simulation of the DQNN algorithm using QuTip [275], a quantum toolbox in Python. The code can be found at [276]. We start with a short description of the code. A welcome side effect is getting a good overview of the parameters and data needed during the algorithm. We conclude the section with numerical results.

#### Algorithm

A description of the code can be found in Algorithm 1. We explained the parts of the algorithm procedure in detail in Section 4.3. Hence we will focus on technical details here only.

The behaviour of the algorithm can be controlled with the parameters step size  $\epsilon$ , learning rate  $\eta$  and number of epochs  $r_T$ . The latter was not discussed in this work until now and determines how often the training algorithm using all data pairs and updating all QNN unitaries is repeated.

Beyond the parameters, the data pairs can be seen as inputs of the algorithm as well. However, the code in [276] includes preparing the training and validation data using a randomly chosen unitary  $Y$ , which is aimed to simulate with the DQNN after successful training.

The output of the algorithm is the trained network, saved in network unitaries and lists of the values of the loss functions, which will be exploited in the following to discuss the training success.

---

**Algorithm 1** Classical simulation of the DQNN algorithm.
 

---

- 1: Set parameters step size  $\epsilon$ , learning rate  $\eta$  and number of epochs  $r_T$
  - 2: Set  $s = 0$
  - 3: Assign the QNN unitaries  $U_j^l(0)$  randomly
  - 4: Provide  $S$  training data pairs  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}$
  - 5: Provide  $N - S$  validation data pairs  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{T}}\rangle\}$
  - 6: **for**  $r_T$  training epochs **do**
  - 7:     Save the values of training loss  $\mathcal{L}_{\text{SV}}$  and validation loss  $\mathcal{L}_{\text{USV}}$
  - 8:     **for** all  $l$  and  $x$  **do**
  - 9:         Feed-forward the input state  $\rho_x^{l-1}(t) = \mathcal{E}_t^{l-1}(\dots \mathcal{E}_t^1(\rho_x^{\text{in}})\dots)$
  - 10:         Back-propagate the output states  $\sigma_x^l(t) = \mathcal{F}_t^{l+1}(\dots \mathcal{F}_t^{L+1}(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|)\dots)$
  - 11:     **end for**
  - 12:     **for** all  $l$  and  $j$  **do**
  - 13:         Calculate the update matrix  $K_j^l(t) = \frac{\eta 2^{m_l-1} i}{S} \sum_x \text{tr}_{\text{rest}}(M_j^l(x, t))$
  - 14:         Update the unitaries  $U_j^l(t + \epsilon) = e^{i\epsilon K_j^l(t)} U_j^l(t)$
  - 15:     **end for**
  - 16: **end for**
  - 17: Feed-forward the input state  $\rho_x^{l-1}(t) = \mathcal{E}_t^{l-1}(\dots \mathcal{E}_t^1(\rho_x^{\text{in}})\dots)$
  - 18: Save the values of training loss  $\mathcal{L}_{\text{SV}}$  and validation loss  $\mathcal{L}_{\text{USV}}$
- 


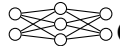


Figure 4.7.: **Training a DQNN.** The validation and training loss converge to the value 1 during the training of a  QNN in  $r_T = 1000$  steps using 10 training pairs and 90 testing pairs (based on a unitary  $Y \in \mathcal{U}(4)$ ),  $\eta = 1$  and  $\epsilon = 0.01$ .

#### 4. Dissipative quantum neural networks

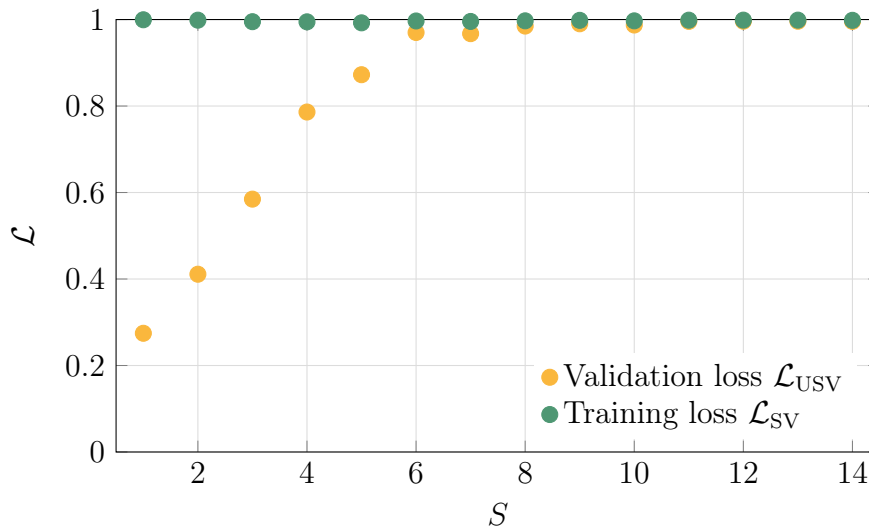
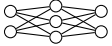


Figure 4.8.: **Generalisation analysis of a DQNN.** This plot shows the training and validation loss after training a  QNN in  $r_T = 1000$  steps with  $\eta = 1$  and  $\epsilon = 0.01$  using 100 data pairs (based on a unitary  $Y$ ), where  $S$  pairs were used for training and  $100 - S$  pairs for testing. We vary the number of  $S$ . All values are averaged over 10 individual training attempts.

In Figure 4.7 the training and validation loss during the whole training algorithm of an exemplary QNN is plotted. The training loss increases faster during the process, but also the validation loss reaches 0.9 after only  $r_T = 363$  epochs. The training and the validation loss end with a value of nearly 1. Since the losses are based on the fidelity of the desired states and the output states of the DQNN, we can conclude that it simulates satisfyingly the unknown unitary  $Y$  on the training input states as well as on the unseen validation input states.

#### Generalisation analysis

In the plot mentioned above  $S = 10$  of 100 data pairs build the training set. The remaining pairs are saved for validation. To study how well the QNN is able to generalise,  $S$  can be varied. Therefore we only compare the end value of the loss function, namely the values obtained in the last line of Algorithm 1. We averaged these values over 10 individual training attempts for all values of  $S \in [1, 14]$  to get a more valid result. Figure 4.8 presents this scenario. As expected, the validation loss increase with the number of training pairs. It is remarkable that using only six of 100 data pairs for training leads to a validation loss over 0.95.

#### Noise analysis

We have already seen a very good generalisation behaviour. Since the approach is to implement the DQNN on currently available quantum devices, we test the

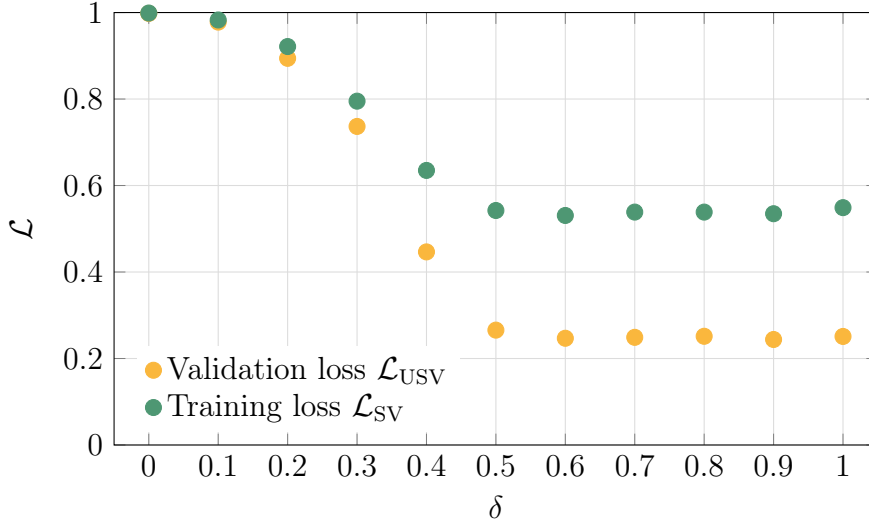
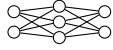


Figure 4.9.: **Testing the noise robustness of a DQNN.** The training and validation loss after training a  QNN changes with the noise parameter  $\delta$ . The training was done in  $r_T = 1000$  steps with  $\eta = 1$  and  $\epsilon = 0.01$  using 20 training and 80 testing pairs (based on a unitary  $Y$ ). The loss values are averaged over 10 individual training attempts.

algorithm for robustness to noise. Therefore we replace the supervised states with  $|\phi_{x,\delta}^{SV}\rangle$ , using a new parameter  $\delta$  and a for every training pair randomly chosen state  $|\psi^{\text{random}}\rangle$  of the same dimension as  $|\phi_x^{SV}\rangle$ , i.e. we build

$$|\phi_{x,\delta}^{SV}\rangle = \frac{(1 - \delta) |\phi_x^{SV}\rangle + \delta |\psi^{\text{random}}\rangle}{\|(1 - \delta) |\phi_x^{SV}\rangle + \delta |\psi^{\text{random}}\rangle\|}.$$

In the same way as when observing the generalisation behaviour in Figure 4.8 we compare the last values of the loss functions and average over 10 completely independent trainings. In Figure 4.9 we plot the validation loss for different values of  $\delta$ .

We can observe that both, the training and the validation loss, take higher values than 0.7 for  $\delta \leq 0.3$ . If the noise exceeds this value of delta, the training loss rests at a value around 0.55 and the validation loss around 0.25.

Discussing only a few examples, we could already observe the great learning behaviour of the DQNN with outstanding generalisation possibility and robustness to noise. More numerical experiments, including 1- and 3-qubit training states and deeper networks, can be found in the appendix in Chapter A.

## 4.6 NISQ device implementation

In the preceding paragraphs, we have already discussed the training results of the DQNN algorithm using a simulation running on a classical computer. However, in

#### 4. Dissipative quantum neural networks

this section, we explain the implementation of the same algorithm on a quantum computer, more precisely on a NISQ device. A short introduction to NISQ devices was given in Section 3.4. Since the implementation presented in Section 4.5 differs from the implementation on these early quantum devices we denote the latter with  $\text{DQNN}_{\text{NISQ}}$ .

#### Implementation of $\text{DQNN}_{\text{NISQ}}$

We already explained the quantum perceptron and the network architecture in detail in Section 4.1. Since the implementation on the NISQ device is based on this, we want to shortly remind the reader of the basic idea at this point. The quantum perceptron is defined as a unitary  $U_j^l$  acting on  $m_{l-1} + 1$  qubits:  $m_{l-1}$  qubits, placed in layer  $l - 1$ , are the input  $\rho^{l-1}$  of the perceptron, and the last qubit belongs to layer  $l$  and is initialised in the zero state. We get the perceptron's output state after applying the unitary  $U_j^l$  to the  $m_{l-1} + 1$  qubits and tracing out the  $m_{l-1}$  input qubits. The perceptron layers of the DQNN can be summarised in unitaries  $U^l = U_{m_l}^l \dots U_1^l$ . We will use this layer notation to describe the implementation of the DQNN on a NISQ device. As before, we assume  $m = m_0 = m_{L+1}$ , hence we train the DQNN to imitate a unitary operation.

As a first and essential step to build a  $\text{DQNN}_{\text{NISQ}}$ , the DQNN's perceptron has to be implemented. In the classical simulation discussed in Section 4.5 the perceptrons were defined by unitary matrices, whose entries would be updated during the training algorithm. As explained in Section 3.5, QNNs can be implemented on a quantum computer via *parametrised quantum circuits* [237, 248–250] consisting of parametrised quantum gates. For this, two aspects need to be considered: the realisation should be universal and the number of gates and parameters small. Here, we present the work of [101], where a balance of these objectives leads to good training results on NISQ device as plotted later in this section. A detailed discussion of over-parametrisation can be found at [277].

To express the perceptron unitaries we apply a result of studies on the implementation of two-qubit gates [278–283]: using a two-qubit canonical gate and twelve single qubit gates, see Figure 4.10, every arbitrary two-qubit unitary can be simulated.




Figure 4.10.: **Implementation of a two-qubit unitary.** An arbitrary two-qubit unitary operation can be implemented using one two-qubit gate and twelve single qubit gates.

Hereby the two-qubit canonical gate  $\text{CAN}(p_x, p_y, p_z)$  is defined as the composition of three two-qubit gates

$$\text{CAN}(p_x, p_y, p_z) = e^{-i\frac{\pi}{2}p_x X \otimes X} e^{-i\frac{\pi}{2}p_y Y \otimes Y} e^{-i\frac{\pi}{2}p_z Z \otimes Z}$$



$$= RXX(p_x\pi) RYY(p_y\pi) RZZ(p_z\pi).$$

using the Pauli matrices  $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ ,  $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ , and  $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$  and the parameters  $p_x, p_y, p_z \in \mathbb{R}$  [282]. It is worth mentioning that the gates  $RXX(p)$ ,  $RYY(p)$  and  $RZZ(p)$  are standard in quantum computing libraries. The single qubit gates  $Y^p$  and  $Z^p$  are parametrised Pauli operators and, up to a phase, equivalent to rotations around the  $y$  and  $z$  axis, namely

$$\begin{aligned} Y^p &\approx R_Y(\pi p) = e^{-i\frac{\pi}{2}pY} \\ Z^p &\approx R_Z(\pi p) = e^{-i\frac{\pi}{2}pZ}. \end{aligned}$$

The sequences of single gates  $Y^p$  and  $Z^p$  with parameters  $p_1, \dots, p_3 \in \mathbb{R}$ , see Figure 4.10, can be rephrased in form of the gate

$$u(p_1, p_2, p_3) = \begin{pmatrix} \cos(p_1/2) & -e^{ip_3} \sin(p_1/2) \\ e^{ip_2} \sin(p_1/2) & e^{i(p_3+p_2)} \cos(p_1/2) \end{pmatrix},$$

which is a standard gate in most of the quantum computing libraries as well. We can proceed in the same way for the parameters  $p_4, \dots, p_{12} \in \mathbb{R}$ .

So far, we have only discussed how to implement two-qubit unitaries. The quantum perceptron unitaries  $U_j^l$  are in general  $m_{l-1} + 1$ -qubit unitaries. These qubit unitaries connect  $m_{l-1}$  qubits from layer  $l - 1$  with one qubit of layer  $l$ . This motivates to replace one of these  $m_{l-1} + 1$ -qubit unitaries with  $m_{l-1}$  two-qubit unitaries acting on one of the qubits in layer  $l - 1$  and the qubit in layer  $l$ . Numerical results presented later in this section justify this idea.

After clarifying the implementation of the quantum perceptron, we can build the network circuit layerwise. Every  $\text{DQNN}_{\text{NISQ}}$  consists of  $M = \sum_{l=0}^{L+1} m_l$  qubits. The first  $m$  qubits are initialised in the input state  $|\phi^{\text{in}}\rangle$ . The remaining qubits are initialised in the zero state  $|0\rangle$ . For every layer, we first apply the single-qubit gates  $u(p_1, p_2, p_3)$  to all qubits of the layer, followed by the unitary  $U^l$ , expressed via  $m_{l-1}$  of the *CAN*-gates. After layer  $l$  is complete the  $m_{l-1}$  input qubits are neglected and the  $m_l$  input qubits act as input qubits for the new layer  $l + 1$ . Different to the classical simulation the feed-forward process is not at the end when the last layer is reached: an extra layer of single-qubit gates follows, see Figure 4.11b. After that, the output of the circuit is a  $m_{L+1}$ -qubit quantum state  $\rho^{\text{out}}$ .

See Figure 4.11 for an example: the implementation of the DQNN depicted in Figure 4.11a is visualised in Figure 4.11b. The first layer unitary  $U^1 = U_3^1 U_2^1 U_1^1$  is written out in two-qubit unitaries (implemented via *CAN*-gates) in Figure 4.11c. The two two-qubit unitaries needed for expressing the perceptron unitary  $U_1^1$  are marked with dashed and dotted lines in Figure 4.11a and Figure 4.11c.

To summarise the description of the implementation we shortly discuss the number of parameters. The  $u$ -gates, placed before every unitary layer and additionally at the end to the output state are described by  $3 \sum_{l=1}^{L+1} m_{l-1} + 3m_{L+1}$  parameters.

#### 4. Dissipative quantum neural networks




Figure 4.11.: **Implementation of a  $\text{DQNN}_{\text{NISQ}}$ .** A  $\text{DQNN}_{\text{NISQ}}$  consisting of two layers of quantum perceptrons and seven qubits (a) can be implemented as quantum circuit using  $u$ -gates and unitary operations representing the layers of the network (b). The first layer including three perceptrons can be decomposed in six two-qubit gates (c). The perceptron  $U_1^1$  is marked exemplarily in dashed and dotted lines (see a and c).

The *CAN*-gates are parametrised via  $3 \sum_{l=1}^{L+1} m_{l-1} m_l$  real numbers. It results that the quantum circuit is described by

$$n^{\text{DQNN}_{\text{NISQ}}} = 3 \sum_{l=1}^{L+1} m_{l-1} (1 + m_l) + 3m_{L+1}$$

parameters.

### Implementation of the training algorithm

As the reader may notice, we so far only described the processing of the input state through the DQNN and becoming the output state. This is only one component, the feed-forward part, of the training algorithm. In the following lines, the rest of the algorithm will be explained.

Analogously to the preceding discussions we use the same learning task, i.e. learn an unknown unitary  $Y \in \mathcal{U}(2^m)$  from a given training set  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}_{x=1}^S$  where  $|\phi_x^{\text{SV}}\rangle = Y |\phi_x^{\text{in}}\rangle$ . Further we use the fidelity as a training loss function, namely

$$\mathcal{L}_{\text{SV}}(\omega_t) = \frac{1}{S} \sum_{x=1}^S \langle \phi_x^{\text{SV}} | \rho_x^{\text{out}}(\omega_t) | \phi_x^{\text{SV}} \rangle,$$

where  $\rho_x^{\text{out}}$  denotes the network's output and the vector  $\omega_t = (\omega_1(t), \dots, \omega_n(t))^T$  with  $n = n^{\text{DQNN}_{\text{NISQ}}}$  comprises the parameters describing the quantum circuit.

In the classical simulation, the training takes place by updating the entries of unitary matrices following an update rule, including an updated matrix. This update matrix was derived using the derivative of the loss function and requires the knowledge of the states of each layer's qubits. Since this is not possible on the NISQ device, the implementation of the algorithm uses parametrised quantum gates. The parameters change during the training in order to maximise the loss function. Gradient descent is a good tool to find out in which direction the change should be done. Note that the gradient descent method can be technically exchanged by any other optimising algorithm.

More precisely the  $n^{\text{DQNN}_{\text{NISQ}}}$  parameters of the quantum circuit are initialised as  $\omega_0$ . After every training epoche all parameters are updated by  $\omega_{t+1} = \omega_t + d\omega_t$ , where  $d\omega_t = \eta \nabla \mathcal{L}_{\text{SV}}(\omega_t)$  using the learning rate  $\eta$  and the gradient is of the form

$$\nabla_k \mathcal{L}_{\text{SV}}(\omega_t) = \frac{\mathcal{L}_{\text{SV}}(\omega_t + \epsilon e_k) - \mathcal{L}_{\text{SV}}(\omega_t - \epsilon e_k)}{2\epsilon} + \mathcal{O}(\epsilon^2).$$

The vectors  $e_k$  are defined as  $e_k^j = \delta_{k,j}$ ,  $k, j = 1, \dots, n^{\text{DQNN}_{\text{NISQ}}}$ . The parameters step size  $\epsilon > 0$  and learning rate  $\eta$  are used analogously to Section 4.5. Note at this point that for calculating the gradient, the training loss function, and therefore the whole quantum circuit, has to be evaluated for the parameters  $\omega_t + \epsilon e_k$  and  $\omega_t - \epsilon e_k$  for every  $k$ .

#### 4. Dissipative quantum neural networks




Figure 4.12.: **Implementation of the of destructive swap test.**  $m$  *CNOT* gates,  $m$  Hadamard gates and  $2m$  measurements are used to calculate the fidelity  $\langle \phi_x^{\text{SV}} | \rho^{\text{out}} | \phi_x^{\text{SV}} \rangle = p \cdot c$ .

The overall procedure of the algorithm resembles with Section 4.1, where the algorithm was presented in detail. The training and validation data, i.e. the set of training pairs and the set of testing pairs, are initialised. The reader is invited to think again of an uncharacterised quantum device acting as a unitary  $Y$  on  $m$ -qubit input states. Further, the step size  $\epsilon$  and learning rate  $\eta$  are determined conformable to the circumstances, and the parameters  $\omega_0$  initialised randomly. As a second step the quantum circuit in Figure 4.11 is executed for all the states in the training set:  $m$  qubits are initialised in a state  $|\phi_x^{\text{SV}}\rangle$  and another  $m$  qubits in  $|\phi_x^{\text{in}}\rangle$ . Note again that it is  $m = m_0 = m_{L+1}$  for learning a unitary. All the other qubits are initialised in the zero state  $|0\rangle$ . The network circuit, depicted in Figure 4.11b, is evaluated.

At this point  $m$  qubits are in the state  $\rho^{\text{out}}$  and another  $m$  qubits in the state  $|\phi_x^{\text{SV}}\rangle$ . We conveniently label the qubits in this order  $q_1, \dots, q_{2m}$ . The next step is to calculate the training loss to update the parameters  $\omega_t$ . For this purpose for all training data pairs, the circuit is executed and the so-called destructive *SWAP*-test [204, 205, 284] is used to execute the fidelity:  $m$  *CNOT* gates and  $m$  Hadamard gates are applied as depicted in Figure 4.12. All qubits are measured in the computational basis and the results are alternating saved in a list, i.e.  $q_1, q_m, q_2, q_{m+1}, \dots, q_{m-1}, q_{2m}$ . If we exemplarily assume  $m = 1$ , i.e. when comparing two one-qubit states, the list can only be of the form  $\{0, 0\}$ ,  $\{0, 1\}$ ,  $\{1, 0\}$  or  $\{1, 1\}$ . Due to the quantum projection noise, we execute these measurements  $M$  times and save the ratio of the occurrence of each list as components of a vector  $p$ . It is reasoned in [205] that the fidelity of the states  $|\phi_x^{\text{SV}}\rangle$  and  $\rho_x^{\text{out}}$  can be obtained by

$$F(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|, \rho_x^{\text{out}}) = p \cdot c,$$

where  $c = (1, 1, 1, -1)^{\otimes m}$ .

Using the *SWAP*-test not only the parameters can be updated via  $\omega_{t+1} = \omega_t + d\omega$ .

Also the validation loss

$$\mathcal{L}_{\text{USV}}(\omega_t) = \frac{1}{N - S} \sum_{x=S+1}^N \langle \phi_x^{\text{USV}} | \rho_x^{\text{out}}(\omega_t) | \phi_x^{\text{USV}} \rangle$$

is calculated in the same way.

In Figure 4.13 the complete quantum circuits, including initialising, the  $\text{DQNN}_{\text{NISQ}}$  algorithm and the *SWAP*-test, is depicted. The box with the title “DQNN” represents the in the above paragraph described DQNN part of the quantum circuit, see Figure 4.11. For a better overview and due to its analogy to the calculation of the training loss, the validation data and the validation process are not depicted.




Figure 4.13.: **Implementation of training a  $\text{DQNN}_{\text{NISQ}}$ .** The training process includes initialising the qubits, performing the  $\text{DQNN}_{\text{NISQ}}$  quantum circuit and executing the *SWAP*-test. The number of used qubits is  $2m+n$ , where  $m = m_0 = m_{L+1}$  and  $n = \sum_{l=1}^L m_l$ .

### Device execution

Before plotting some results, we shortly have to mention a third loss function, next to the training and validation loss. The *identity loss* is worked out using  $Y = \mathbb{1}$  and parameters which make the network, assuming there would be no noise, act as the identity. Nevertheless, indeed, all the gates are still applied and add noise to the circuit. Hence this loss gives a good insight for the best possible training loss an ideally trained network could generate.

In Figure 4.14 the training of an exemplary  $\text{DQNN}_{\text{NISQ}}$  learning an unknown unitary  $Y \in \mathcal{U}(4)$  is depicted. Based on the quantum no free lunch theorem, see Chapter 5, four training pairs are used to achieve the depicted training success. Another four data pairs were used to calculate the validation loss.

Before the training algorithm was executed, the parameters were initialised in the range  $[0, 2\pi)$ . The algorithm was carried out on the 7-qubit quantum device *ibmq\_casablanca* hosted by IBM [59]. For the implementation and execution of the algorithm, we used the open-source SDK Qiskit [285], which allows the decomposition of the above-explained parametrised quantum circuits to a circuit able to execute by the IBM quantum devices. It is noteworthy at this point to

#### 4. Dissipative quantum neural networks




Figure 4.14.: **Training a DQNN<sub>NISQ</sub>** . A  $\otimes$  DQNN<sub>NISQ</sub> network has been trained on *ibmq\_casablanca* with  $\epsilon = 0.5$  and  $\eta = 1.0$  in 100 epochs. The training loss is computed using 4 training pairs (based on a unitary  $Y \in \mathcal{U}(4)$ ). After every fifth epoch, the identity loss using 4 output training states and the validation loss is additionally measured using 4 validation data pairs, are analysed as well.

underline again that the training took place in a hybrid manner, meaning that the loss functions were evaluated indeed by the quantum implementation however the update of the parameters happened classically.

It looks like the training and validation loss are correlated to the variation of the identity loss. This strengthens the assumption that the identity loss arises from the noise of the quantum device and is not just based on statistical errors. We can understand from the plot that the training loss exceeds the identity loss after a few training epochs and come to the result that the network is able to generalise the information provided through the training data despite the high noise levels. Because of the small number of parameters, effects of barren plateaus are not encountered.

### 4.7 Comparison to quantum approximate optimisation algorithm

We want to close this chapter with a comparison of the above comprehensively discussed DQNN<sub>NISQ</sub> with another QNN architecture, the *quantum approximate optimisation algorithm* (QAOA) [103–105]. For this, we follow the work of [101] where not only both QNNs were challenged to learn an unknown unitary, but also the noise tolerance of the two approaches are compared, when implemented on an IBM NISQ device [59]. The QAOA leads not only to solutions of combinatorial problems [235, 256, 286–292] and is universal for quantum computation [293], but

#### 4.7. Comparison to quantum approximate optimisation algorithm

was also successfully used to learn unknown unitaries [272]. Since the latter was the leading learning example in this chapter, a comparison of the  $\text{DQNN}_{\text{NISQ}}$  and QAOA suggests itself. We begin with a short introduction to the QAOA. A comparison with numerical results ensues.

##### Quantum approximate optimisation algorithm

The QAOA can be interpreted in various ways. In the following, we discuss the QAOA according to [103]. There, this algorithm is described as a QNN and this suits the comparison to the  $\text{DQNN}_{\text{NISQ}}$  the best. Despite this similar implementation, these two approaches differ in crucial aspect: whereas the  $\text{DQNN}_{\text{NISQ}}$  works in a dissipative manner and a perceptron acts on layers of different qubits, in the QAOA setting, a perceptron is defined as a sequence of operations, and all perceptrons act on the identical qubits.

The QAOA is sometimes also called *quantum alternating operator ansatz*, due to the fact that it is working with a sequence of alternating unitary operations  $e^{-iAp_l}, e^{-iBk_l} \in \mathcal{U}(d)$ , where  $p_l, k_l \in \mathbb{R}$ . The matrices  $A$  and  $B$  are hermitian and initialised using the Gaussian unitary ensemble [103].  $\tau$  of these operator pairs can be phrased as

$$\mathcal{U} = e^{-iBk_\tau} e^{-iAp_\tau} \dots e^{-iBk_1} e^{-iAp_1}.$$

The output state of the algorithm can be written as  $\rho^{\text{out}} = \mathcal{U} |\phi^{\text{in}}\rangle \langle \phi^{\text{in}}| \mathcal{U}^\dagger$ .

Such an operator sequence can be interpreted as a  $(\tau + 1)$ -layer QNN, where the number of neurons per layer  $m$  is equal in every layer. A layer  $l$  is defined as an operator pair  $e^{-iAp_l}, e^{-iBk_l}$  acting on all  $m$  neurons of layer  $l - 1$ . The changed neurons are the inputs for the layer  $l + 1$ . This procedure is depicted for two neurons  $q_1$  and  $q_2$  in Figure 4.15a.

The implementation of QAOA as a quantum circuit is very simple. The  $m$  qubits are initialised in the input state  $|\phi^{\text{in}}\rangle$  with dimension  $d = 2^m$ . The number of sequences  $\tau$  is chosen to be  $d^2/2$ . In that way the QAOA leads to an optimal solution [272] and the number of parameters is  $n^{\text{QAOA}} = d^2 = 4^m$ . See Figure 4.15b for the illustration of the implementation of an exemplary QAOA circuit.




Figure 4.16.: **Implementation of training the QAOA.** After the qubits are initialised the QAOA quantum circuit can be executed. The *SWAP*-test allows calculating the values of the loss functions. The number of used qubits is  $2m$ .

The QAOA circuit can be trained equivalently to the training of the

#### 4. Dissipative quantum neural networks




Figure 4.15.: **Implementation of the QAOA.** A two-qubit QAOA can be represented as a  $(\tau + 1)$ -layer QNN, where the same qubits are used in every layer (a). The implementation as a quantum circuit consists of a sequence of alternating unitary operations (b).




Figure 4.17.: **Training a QAOA.** The plot shows the training of an  $m = 2$  QAOA on *ibmq\_casablanca* with  $\epsilon = 0.15$  and  $\eta = 1.0$  in 100 epochs. The training loss is computed using 4 training pairs. After every fifth epoch, the identity loss using 4 output training states and the validation loss is additionally measured using 4 validation data pairs, are analysed as well.



#### 4.7. Comparison to quantum approximate optimisation algorithm

DQNN<sub>NISQ</sub> circuit including the following steps: The training and validation data set consisting of pairs of states  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}_{x=1}^N$  are prepared, where  $|\phi_x^{\text{SV}}\rangle = Y |\phi_x^{\text{in}}\rangle$ . The circuit parameters are initialised randomly. The input states are carried through the quantum circuit and the resulting output states and the supervised states are used to calculate the validation with the *SWAP*-test. Gradient descent is performed to update the parameters. Despite the equivalence to the DQNN<sub>NISQ</sub> training, the training algorithm of a QAOA is depicted as a circuit in Figure 4.16.

In Figure 4.17 the training of an  $m = 2$  QAOA is depicted while learning an unknown unitary  $Y \in \mathcal{U}(4)$ . Four training pairs are used to achieve the depicted training success. Another four data pairs were used to calculate the validation loss. The identity loss is defined as explained in Section 4.6. In analogy to Figure 4.14 the QAOA was carried out on the 7-qubit quantum device *ibmq\_casablanca* by IBM [59].

In comparison to Figure 4.14, where the same device was used to train the DQNN<sub>NISQ</sub>, the validation loss does not exceed the identity loss in the process of training the QAOA. However, it seems to be correlated to the identity loss. Since comparing two results of single training sessions is not very fair, we cite in the following the results on generalisation and gate noise analysis of [101].

#### Comparison of DQNN<sub>NISQ</sub> and QAOA

In the following, we compare both networks, the DQNN<sub>NISQ</sub> and the QAOA, with  $m = 2$  input and output qubits. The DQNN<sub>NISQ</sub> algorithm comes in the form of four qubits, two of them in both of the layers. The QAOA includes only two qubits, and we choose  $\tau = 8$ . To calculate the loss functions, two more qubits for training each of the QNNs are needed.

For the training  $S = 4$  pairs of training data,  $N - S = 4$  pairs of validation data are used. Additional four state pairs are used to evaluate the identity loss.

A difference is, that the  $n^{\text{DQNN}_{\text{NISQ}}} = 24$  parameters of the DQNN<sub>NISQ</sub> are initialised in the range  $[0, 2\pi)$ , where the  $n^{\text{QAOA}} = 16$  parameters needed for the QAOA are initialised in the range  $[-1, 1]$ . Further, different parameters  $\eta$  and  $\epsilon$  has been found as optimal (see the captions of Figure 4.18 and Figure 4.19 for examples).

As already discussed in Section 4.2 we are interested in the networks generalisation capabilities and use the validation loss  $\mathcal{L}_{\text{USV}}$  for studying it. Both network's quantum circuits were executed on a simulator of *ibmq\_casablanca* imitating the real-time noise of this NISQ device. After repeating the same training with the same circumstances and the same number of supervised pairs  $S$  but different initial parameters, we get an average of the testing loss for each  $S$ . The results of this numerical experiment can be found in Figure 4.18.

First of all, we can say that both, the DQNN<sub>NISQ</sub> and the QAOA, can generalise the information given through the training data pairs. As expected, the validation loss gets larger with the increasing number of training pairs. It also becomes clear that the DQNN<sub>NISQ</sub> reaches higher values which have to be observed together with

#### 4. Dissipative quantum neural networks




Figure 4.18.: **Generalisation analysis.** This figure compares the generalisation behaviours of a  $\otimes$  DQNN<sub>NISQ</sub> ( $\epsilon = 0.5$ ,  $\eta = 1.0$ ) and a  $m = 2$  QAOA ( $\epsilon = 0.15$ ,  $\eta = 0.1$ ), each trained with  $S$  training pairs for  $r_T = 1000$  epochs on a *ibmq\_casablanca* simulator.

the higher identity loss. Since the training results seem to be strongly affected by the noise of the executing quantum device, a noise analysis of both of the QNN attempts will be discussed in the succeeding paragraph.

Two primary types of noise are the readout noise, occurring during the measurements, and the gate noise [294]. Since both networks are trained, read out and updated with the same methods but consist of different gates we focus only on the gate noise in the following.

Using a depolarising quantum error channel [52] we are able to test the influence of gate noise on the QNNs. In the numerical experiments, the channel was parametrised by the depolarisation probabilities  $\lambda^g = k\lambda_0^g$  for the basis gates  $g = \text{CNOT}, \text{SX}, \text{RZ}$  and scaling factor  $k$ . To simulate the gate errors of a special NISQ device, the parameter  $\lambda_0^g$  has to be chosen appropriately, for example  $\lambda_0^{\text{CNOT}} = 3.14 \times 10^{-2}$ ,  $\lambda_0^{\text{SX}} = 1.18 \times 10^{-3}$  and  $\lambda_0^{\text{RZ}} = 0$  for the approximation of *ibmq\_16\_melbourne* [295].

In Figure 4.19 the values of identity, validation and training losses of the DQNN<sub>NISQ</sub> and QAOA for different error probability factors  $k$  are shown. To value the results, the reader must consider that the noise factor  $k = 1$  corresponds to the gate noise of currently available NISQ devices. The results are gained using the of the qubit coupling map IBM device *ibmq\_16\_melbourne*.

Both attempts give excellent results, i.e. all losses equal one, for the noise-free case  $k = 0$ . When increasing the gate noise, the DQNN<sub>NISQ</sub> leads to a higher

#### 4.7. Comparison to quantum approximate optimisation algorithm




Figure 4.19.: **Gate noise analysis.** Here, the robustness to gate noise of a  $\otimes$  DQNN<sub>NISQ</sub> ( $\epsilon = 0.25$ ,  $\eta = 0.5$ ) and a  $m = 2$  QAOA ( $\epsilon = 0.05$ ,  $\eta = 0.05$ ) is compared. Each are trained with  $S$  training pairs for  $r_T = 1000$  epochs on a *ibmq\_16\_melbourne* simulator.

identity loss compared to the QAOA. This concedes in a higher training and validation loss.

To conclude, we can say that both QNNs succeed in learning an unknown unitary operation. Albeit it also becomes clear that the DQNN<sub>NISQ</sub> algorithm attends this learning task more reliably than the QAOA and is less susceptible to gate noise. The results indicate that the DQNN<sub>NISQ</sub> is more suitable for learning an unknown unitary operation on a NISQ device of the current stage compared to the QAOA. Nevertheless, it needs to be said that the noise still forbids reaching high values in the loss functions in both algorithms.

In the preceding sections, we studied the behaviour of the DQNN algorithm characterising a unitary operation, using numerical results gained with classical simulation or the execution on a NISQ device, where we could see that the DQNNs can be successfully trained and perform well on unseen data. In contrast, the following chapter discusses the theoretical boundaries of such a learning process and allows observing the performance of the DQNN from a new perspective.



# 5

---

## No free lunch theorem

---

Due to the continuous rapid progress in the field of quantum learning theory, it is also significant to understand the ultimate limits for quantum learning devices and methods. The preceding chapter introduced a *quantum neural network* (QNN) structure that can successfully learn an unknown unitary  $Y$  even with limited data. Before we continue in Chapter 6 and Chapter 7 with discussing some further opportunities these QNNs offer, we want to take a step back and look at the boundaries such a learning process has imposed by the *quantum no free lunch* (QNFL) theorem [100].

Therefore we assume the studied device can be modelled as a unitary process and is trained with quantum examples. We will describe in the following how to find an optimal lower bound on the probability that such a QNN gives an incorrect output for a random input. This bound equips us with a helpful metric to check how well particular *quantum machine learning* (QML) architectures and algorithms perform.

The here discussed result can be classed within the research on fundamental information-theoretic limits on quantum learning [70, 72, 75, 271, 296, 297] and is related to the work on the optimal quantum learning of unitary operations considering storage and later retrieval of unknown quantum processes. The version of the QNFL theorem we discuss here was stated in [100] and assumes the usages of training data consisting of quantum states for the goal of learning an unknown unitary process. This work was later generalised by [298] to the case where these quantum states can be entangled to a reference system.

Since the here presented bound is based on the quantum version of the *no free lunch* (NFL) theorem [299], a result of classical learning theory, we begin with a brief explanation of the latter in Section 5.1. In Section 5.2, we derive the

## 5. No free lunch theorem

QNFL theorem following [100]. During the derivation, some Haar measure integral identities for the unitary group are used. We present the proofs of these collectively after the derivation for a neat overview. The derived bound is then used to check the in Chapter 4 presented *dissipative quantum neural network* (DQNN) algorithm. We end the chapter with a short but essential note on orthonormal training pairs.

### 5.1 Classical no free lunch theorem

The classical NFL theorem states that an optimisation algorithm that performs better for one class of problems must perform worse for another class. We simply get no class "for free". Despite different mathematical formulations of this theorem existing, we follow the notation of [274] hereafter.

After introducing the principle of supervised ML in Chapter 2, the following setting should be very familiar to the reader: let  $X$  and  $Y$  be two finite sets. We name  $X$  the *input* and  $Y$  the *output set*. For an unknown function  $f : X \rightarrow Y$  we call the function  $h : X \rightarrow Y$  a *hypothesis*. This hypothesis is based on a subset  $\mathcal{S} \subset X \times Y$  containing training pairs  $\mathcal{S} = \{(x_x, f(x_x)) \mid x_x \in X, f(x_x) \in Y, j = 1, 2, \dots, S\}$ , where  $f(x_x)$  is the desired output given input  $x_x$ . It should yield  $h_S(x_x) = f(x_x)$  for all  $j = 1, 2, \dots, S$ . Note that we assume  $S < |X|$ , otherwise there would be nothing to predict.

The big question leading to the NFL theorem is: how well does a given hypothesis perform? To determine this we define the *risk*  $R_f(h)$  as the probability that  $h_S$  gives the wrong answer, namely

$$R_f(h) \equiv \mathbb{P}[h_S(x) \neq f(x)].$$

When averaging over all possible training sets with  $S$  elements ( $\mathbb{E}_{\mathcal{S}}$ ) and over all possible functions from  $X$  to  $Y$  ( $\mathbb{E}_f$ ), the NFL theorem takes the shape

$$\mathbb{E}_f [\mathbb{E}_{\mathcal{S}} [R_f(h_S)]] \geq \left( \frac{|X| - S}{|X|} \right) \left( \frac{|Y| - 1}{|Y|} \right) = \left( 1 - \frac{S}{|X|} \right) \left( 1 - \frac{1}{|Y|} \right), \quad (5.1)$$

where  $h_S$  denotes an optimal hypothesis given the training data  $\mathcal{S}$ . It follows naturally that if a learning algorithm performs better at predicting  $f$  for one part of the problem, there are other parts where the algorithm will perform worse.

Inspecting Equation (5.1) in detail, we can understand the classical NFL theorem in the following way: if  $h_S$  is perfectly trained for  $S$  data pairs, getting the wrong answer can be explained in two steps. First of all,  $x \notin \mathcal{S}$ , otherwise the optimal hypotheses would give us the correct answer. This circumstance is reflected in the first factor. The second factor describes the probability that the hypothesis guesses any point  $y \in Y$  but the right one.

### No free lunch theorem for invertible functions

When comparing the classical and the, in the next section introduced, QNFL theorem, we should bear in mind a crucial point: a classical function  $f$  does not have to be invertible. If we have  $f$  determined on some subset of the inputs, we have no information about the action of  $f$  on the complement of this subset. In contrast, unitary quantum operations are always invertible. Thus if we already have the operation determined on one subspace, we know that the operation takes the complementary subspace to the complementary subspace of the output.

Therefore in [100] it is argued toward comparing the QNFL theorem with a *classical NFL theorem for invertible functions*. This bound is derived as

$$\mathbb{E}_f [\mathbb{E}_{\mathcal{S}} [R_f(h_{\mathcal{S}})]] \geq \left( \frac{|X| - S}{|X|} \right) \left( \frac{|X| - S - 1}{|X|} \right) = 1 - \frac{S + 1}{|X|}.$$

The above-gained intuition of Equation (5.1) helps us to understand this new bound quickly: if our perfectly trained hypothesis is invertible, it is in particular injective assuming  $|X| = |Y|$ . It follows that we have fewer chances of being wrong when guessing the correct point in  $Y$  because  $S$  of these training points are already omitted: we get the classical NFL theorem for invertible functions if we substitute the cardinality of  $Y$  by  $|X| - S$  in the numerator of the second factor, where  $S$  is the number of data pairs used for training the hypothesis.

## 5.2 Quantum no free lunch theorem

After gaining a good intuition of the classical NFL theorem, we will derive the quantum analogue in the coming paragraphs following the work of [100].

For this purpose we replace the sets  $X$  and  $Y$  mentioned in the last section with an *input Hilbert space*  $\mathcal{H}_{\text{in}}$  and *output Hilbert space*  $\mathcal{H}_{\text{out}}$ . To replace the role of the expressions  $|X|$  and  $|Y|$  we use  $d = \dim(\mathcal{H}_{\text{in}})$  and  $d' = \dim(\mathcal{H}_{\text{out}})$ .

Moreover, the aim is not to define an unexplored function  $f$  but to determine an unknown unitary operation  $Y$ . As we are familiar with from Chapter 4 the training data is of the form  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}$ ,  $x = 1, 2, \dots, N$ , where  $|\phi_x^{\text{in}}\rangle \in \mathcal{H}_{\text{in}}$  and  $|\phi_x^{\text{SV}}\rangle = Y |\phi_x^{\text{in}}\rangle \in \mathcal{H}_{\text{out}}$  and we formulate the hypothesis as a unitary  $U$ . We assume that the unitary  $U$  can exactly reproduce the action of the unknown unitary  $Y$  on the training pairs of the set  $\mathcal{S}$ , namely

$$U |\phi_x^{\text{in}}\rangle = |\phi_x^{\text{SV}}\rangle = Y |\phi_x^{\text{in}}\rangle, \quad x = 1, 2, \dots, N.$$

The most interesting part of generalising the classical NFL theorem is to find a good *quantum risk* analogously to the classical risk in Section 5.1. Here we choose the square of the trace norm distance  $\|A\|_1 \equiv \frac{1}{2} \text{tr} |A|$  between the outputs of  $Y$  and  $U$  applied to the same input, averaged over all pure states. A discussion of the risk in the quantum setting can be found in [75], details concerning the trace

## 5. No free lunch theorem

norm in [52]. Hence we formulate the risk as

$$\begin{aligned}
R_Y(U) &\equiv \int d|\psi\rangle \|Y|\psi\rangle\langle\psi|Y^\dagger - U|\psi\rangle\langle\psi|U^\dagger\|_1^2 \\
&= 1 - \int d|\psi\rangle |\langle\psi|Y^\dagger U|\psi\rangle|^2 \\
&= 1 - \frac{1}{d(d+1)} (d + |\text{tr}(Y^\dagger U)|^2) \\
&= \frac{d}{d+1} - \frac{1}{d(d+1)} |\text{tr}(Y^\dagger U)|^2,
\end{aligned} \tag{5.2}$$

where the integral runs over pure states induced by a Haar-measure-distributed unitary  $W$  applied to an arbitrary state  $d|\psi\rangle \equiv dW|0\rangle$  [209, 210]. A short introduction to the Haar-measure can be found Section 3.2. We explicitly evaluate the integral  $S_4 \equiv \int dY Y^\dagger \otimes Y^\dagger \otimes Y \otimes Y$  and thus the identity used above  $\int d|\psi\rangle |\langle\psi|Y^\dagger U|\psi\rangle|^2 = -\frac{1}{d(d+1)} (d + |\text{tr}(Y^\dagger U)|^2)$  in Proposition 5.2. In the same way as in the classical case the risk defines the probability that the hypothesis fails to reproduce the action of  $Y$  or in other words the probability of incorrectly learning a unitary process.

The next step is to average the risk  $R_Y(U)$  over all possible training sets  $\mathcal{S}$  and unitaries  $Y$ . The first average is trivial. When averaging over all unitaries  $Y$ , we need to have in mind that  $U$  is the best guess for our unknown unitary operation  $Y$  given the training set  $\mathcal{S}$  and therefore  $U$  depends on  $Y$  in a not obvious way. Consequently we formulate

$$\int dY R_Y(U) = \frac{d}{d+1} - \frac{1}{d(d+1)} \int dY |\text{tr}(Y^\dagger U)|^2.$$

To evaluate the integral, we use the fact that  $U$  and  $Y$  act on the training set  $\mathcal{S}$  in the same way. Due to linearity we know further that the unitaries  $Y$  and  $U$  are equal on the subspace  $\mathcal{H}_S \equiv \text{span}(\mathcal{S})$ .

Let  $\mathcal{H}_S^\perp$  be the subspace complementary to  $\mathcal{H}_S$ . Although we have no insight into the action of  $Y$  on this space, we can use the properties  $Y$  has as a unitary. To understand these constraints, we take a closer look at the unitary  $Y^\dagger U$ . Due to the properties of the training set and the direct sum decomposition  $\mathcal{H}_{\text{in}} = \mathcal{H}_S \oplus \mathcal{H}_S^\perp$  we can write  $Y^\dagger U$  in the following block decomposition:

$$Y^\dagger U = \left( \begin{array}{c|c} \mathbb{1}_n & A \\ \hline B & W \end{array} \right) = \left( \begin{array}{c|c} \mathbb{1}_n & \mathbf{0} \\ \hline \mathbf{0} & W \end{array} \right) = \mathbb{1}_n \oplus W.$$

$\mathbb{1}_n$  denotes the  $S$ -dimensional identity on the subspace  $\mathcal{H}_S$ , and  $A$ ,  $B$ , and  $W$  stand for  $S \times (d-S)$ ,  $(d-S) \times S$ , and  $(d-S) \times (d-S)$  block matrices. For the simple reason that  $Y^\dagger U$  is a unitary, the norm of each row and column vector has to be equal to 1. It follows that  $A = B = 0$  and  $W$  is unitary. Because of this



form, the trace of  $Y^\dagger U$  can be written as a sum of traces over  $\mathcal{H}_S$  and  $\mathcal{H}_S^\perp$ , namely

$$\begin{aligned} |\operatorname{tr}(Y^\dagger U)|^2 &= |\operatorname{tr}_{\mathcal{H}_S}(Y^\dagger U) + \operatorname{tr}_{\mathcal{H}_S^\perp}(Y^\dagger U)|^2 \\ &= |n + \operatorname{tr}_{\mathcal{H}_S^\perp}(Y^\dagger U)|^2 \\ &= S^2 + 2S\operatorname{Re}(\operatorname{tr}_{\mathcal{H}_S^\perp}(Y^\dagger U)) + |\operatorname{tr}_{\mathcal{H}_S^\perp}(Y^\dagger U)|^2 \\ &= S^2 + 2S\operatorname{Re}(\operatorname{tr}(W)) + |\operatorname{tr}(W)|^2. \end{aligned}$$

We have no further constraints on the unitary  $W$  and need to guess  $W$  randomly with respect to Haar measure on the unitary group  $\mathcal{U}(d - S)$ . It follows that the average over  $Y$  is cut down to the average of  $W$  over the unitary group  $\mathcal{U}(d - S)$ . As a result we get

$$\int dY |\operatorname{tr}(Y^\dagger U)|^2 = \int dW (n^2 + 2S\operatorname{Re}(\operatorname{tr}(W)) + |\operatorname{tr}(W)|^2). \quad (5.3)$$

The second integrand on the right hand side is linear in  $W$  and vanishes. To demonstrate this one can simply substitute  $W' = \exp^{-i\theta} dW$  and get

$$\begin{aligned} I &= \int dW \operatorname{tr}(W) \\ &= \int dW' \exp^{-i\theta} \exp^{-i\theta} \operatorname{tr}(W') \\ &= \int dW' \exp^{-2i\theta} \operatorname{tr}(W') \\ &= \exp^{-2i\theta} I. \end{aligned}$$

It follows  $I = 0$ . The last summand in Equation (5.3) equals  $\int dY |\operatorname{tr}(Y)|^2 \equiv \operatorname{tr}(S_2) = \frac{1}{d} \operatorname{tr}(SWAP) = 1$ . The identity  $S_2 = \frac{1}{d} SWAP$  is shown in Proposition 5.1 and the integral results in

$$\int dY |\operatorname{tr}(Y^\dagger U)|^2 = S^2 + 1.$$

Overall the QNFL theorem is of the form

$$\mathbb{E}_Y[\mathbb{E}_S[R_Y(U)]] \geq 1 - \frac{1}{d(d+1)}(S^2 + d + 1).$$

### Haar measure itegral identities for the unitary group

Above, we described the derivation of a quantum analogue to the classical NFL theorem. In this context we evaluated integrals running over pure states induced by a Haar-measure-distributed unitary  $W$  applied to an arbitrary state  $d|\psi\rangle \equiv dW|0\rangle$  [209, 210]. More precisely, we need two results for the derivation, in the following formulated and proven as Proposition 5.1 and Proposition 5.2.

As already introduced in Section 3.2, we write the integral over the unitary

## 5. No free lunch theorem

group  $\mathcal{U}(d)$  of  $d \times d$  matrices of  $f(Y)$  with respect to Haar measure as

$$I = \int dY f(Y),$$

where  $f(Y)$  is a matrix-valued function on  $\mathcal{U}(d)$ . Recall that The Haar measure is left- and right-invariant with respect to shifts via multiplication. With this in mind we show that  $S_2$  can be expressed using the *SWAP* operation. Note that we use tensor network diagrams for the derivations. An introduction to the topic can be find at [300].

**Proposition 5.1.**

$$S_2 \equiv \int dY Y^\dagger \otimes Y = \frac{1}{d} \text{SWAP}$$

*Proof.* To proof this identity we note that  $S_2$  fulfils

$$S_2 = (\mathbb{1} \otimes e^{i\epsilon X}) S_2 (e^{-i\epsilon X} \otimes \mathbb{1})$$

for any hermitian operator  $X$  and  $\epsilon > 0$  infinitesimally small. Expanding the right hand side to the first order in  $\epsilon$  leads to

$$0 = i\epsilon(\mathbb{1} \otimes X) S_2 - i\epsilon S_2 (X \otimes \mathbb{1}),$$

and further to

$$S_2 (X \otimes \mathbb{1}) = (\mathbb{1} \otimes X) S_2.$$

The expression is valid for any hermitian operator. Therefore we can assume in the following that  $X$  is each of a Hilbert-Schmidt orthonormal hermitian operator basis  $\lambda^\alpha$ ,  $\alpha = 0, 1, \dots, d^2 - 1$ , with  $\text{tr}(\lambda^\alpha \lambda^\beta) = \delta^{\alpha\beta}$ . If we replace  $X$  with  $\lambda^\alpha$ , multiply the expression on the right by  $\lambda^\alpha \otimes \mathbb{1}$  and sum over  $\alpha$  we end up with the equation

$$\sum_{\alpha} S_2(\lambda^\alpha \lambda^\alpha \otimes \mathbb{1}) = \sum_{\alpha} (\mathbb{1} \otimes \lambda^\alpha) S_2(\lambda^\alpha \otimes \mathbb{1}). \quad (5.4)$$

We can express the identity  $\text{SWAP} = \sum_{\alpha} \lambda^\alpha \otimes \lambda^\alpha$  in such a diagram as

$$\text{Diagram of two crossing lines} = \sum_{\alpha} \begin{array}{c} \text{---} \circ \lambda^\alpha \text{---} \\ \text{---} \circ \lambda^\alpha \text{---} \end{array}$$

When wiring together the outputs of this diagram we end up with

$$\sum_{\alpha} \text{---} \circ \lambda^\alpha \text{---} \circ \lambda^\alpha \text{---} = \sum_{\alpha} \begin{array}{c} \text{---} \circ \lambda^\alpha \text{---} \\ \text{---} \circ \lambda^\alpha \text{---} \end{array} = \text{Diagram of two lines crossing and looping back} = d \mathbb{1}$$

Equation (5.4) can be represented as

$$\sum_{\alpha} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \begin{array}{c} \lambda^{\alpha} \\ \lambda^{\alpha} \end{array} = \sum_{\alpha} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \begin{array}{c} \lambda^{\alpha} \\ \lambda^{\alpha} \end{array}$$

Using the above described identities leads to

$$d \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} = \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array}$$

Now we can use the integral representation of  $S_2$ , namely

$$\begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} = \int dY \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array}$$

and connect the outputs. In tensor network diagrams this can be depicted as

$$\begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} = \int dY \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} = \text{---}$$

Finally we end up with

$$\begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} = \frac{1}{d} \begin{array}{c} \text{---} \\ \text{---} \end{array}$$

which is another way to denote  $\int dY Y^{\dagger} \otimes Y = \frac{1}{d} \text{SWAP}$ . □

Using network diagram also the second needed identity used in the proof of the QNFL theorem can be proven.

**Proposition 5.2.**

$$\begin{aligned} S_4 &\equiv \int dY Y^{\dagger} \otimes Y^{\dagger} \otimes Y \otimes Y \\ &= \frac{1}{d^2 - 1} \begin{array}{c} \text{---} \\ \text{---} \end{array} - \frac{1}{d(d^2 - 1)} \begin{array}{c} \text{---} \\ \text{---} \end{array} + \frac{1}{d^2 - 1} \begin{array}{c} \text{---} \\ \text{---} \end{array} - \frac{1}{d(d^2 - 1)} \begin{array}{c} \text{---} \\ \text{---} \end{array} \end{aligned}$$

*Proof.* The tensor-network diagram of the stated identity is

5. No free lunch theorem

$$\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \text{---} S_4 \text{---} = \int dY \begin{array}{c} \text{---} \text{---} Y^\dagger \text{---} \\ \text{---} \text{---} Y^\dagger \text{---} \\ \text{---} \text{---} Y \text{---} \\ \text{---} \text{---} Y \text{---} \end{array}$$

In analogy to the proof of Proposition 5.1 we get the equation

$$S_4(X \otimes \mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1}) + S_4(\mathbb{1} \otimes X \otimes \mathbb{1} \otimes \mathbb{1}) = (\mathbb{1} \otimes \mathbb{1} \otimes X \otimes \mathbb{1})S_4 + (\mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1} \otimes X)S_4$$

by mapping  $Y$  to  $e^{i\epsilon X}Y$  with infinitesimal small  $\epsilon$ . Further we follow the same proof strategy and set  $X = \lambda_\alpha$ , multiply on the right by  $\lambda_\alpha \otimes \mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1}$ , and sum over  $\alpha$ . We end up with

$$\sum_\alpha \begin{array}{c} \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \\ \text{---} \end{array} S_4 \text{---} + \begin{array}{c} \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \\ \text{---} \end{array} S_4 \text{---} = \sum_\alpha \begin{array}{c} \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \\ \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \end{array} S_4 \text{---} + \begin{array}{c} \text{---} \text{---} \lambda^\alpha \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \text{---} \lambda^\alpha \text{---} \end{array} S_4 \text{---}$$

Defining an operator  $M$  as

$$\begin{array}{c} \text{---} \\ \text{---} \end{array} \text{---} M \text{---} = \text{---} \text{---} + \frac{1}{d} \begin{array}{c} \text{---} \\ \text{---} \end{array} \text{---}$$

and using the identities derived in the proof of Proposition 5.1 leads to

$$d \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \text{---} S_4 \text{---} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} M \text{---} = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \text{---} S_4 \text{---} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} + \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \text{---} S_4 \text{---} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array}$$

For the first summand we can use the identity

$$\int dY \begin{array}{c} \text{---} \text{---} Y^\dagger \text{---} \\ \text{---} \text{---} Y^\dagger \text{---} \\ \text{---} \text{---} Y \text{---} \\ \text{---} \text{---} Y \text{---} \end{array} = \begin{array}{c} \text{---} \\ \text{---} \end{array} \text{---} S_2 \text{---} = \frac{1}{d} \begin{array}{c} \text{---} \\ \text{---} \end{array}$$

The second summand can be rewritten similar, and we end up with

5.2. Quantum no free lunch theorem

$$d \quad S_4 \quad M = \frac{1}{d} \text{crossing} + \frac{1}{d} \text{crossing with loop}$$

By assuming the inverse is as well a linear combination of the identity and the SWAP operation, and using  $\text{SWAP} \times \text{SWAP} = \mathbb{1}$  we can show that the inverse of  $M$  is given by

$$M^{-1} = \frac{d^2}{d^2 - 1} \text{line} - \frac{d}{d^2 - 1} \text{crossing}$$

It remains to multiply both sides by  $M^{-1} \otimes \mathbb{1} \otimes \mathbb{1}$  and the result is the aimed identity

$$S_4 = \frac{1}{d^2 - 1} \text{crossing} - \frac{1}{d(d^2 - 1)} \text{crossing with loop} + \frac{1}{d^2 - 1} \text{crossing with loop} - \frac{1}{d(d^2 - 1)} \text{crossing with loop}$$

□

For the sake of completeness we shortly note how this representation of  $S_4$  helps to derive Equation (5.2) using  $X \equiv Y^\dagger U$ , namely

$$\int d|\psi\rangle |\langle\psi|X|\psi\rangle|^2 = \int dY \langle 0|Y^\dagger X^\dagger Y|0\rangle \langle 0|Y^\dagger XY|0\rangle$$

$$= \int dY \left[ \text{diagram with four lines and two loops} \right]$$

$$= \frac{|\text{tr}(X)|}{d^2 + 1} - \frac{d}{d(d + 1)} + \frac{d}{d^2 + 1} - \frac{|\text{tr}(X)|}{d(d + 1)}$$

$$= -\frac{1}{d(d + 1)} (d + |\text{tr}(X)|^2).$$

## 5. No free lunch theorem

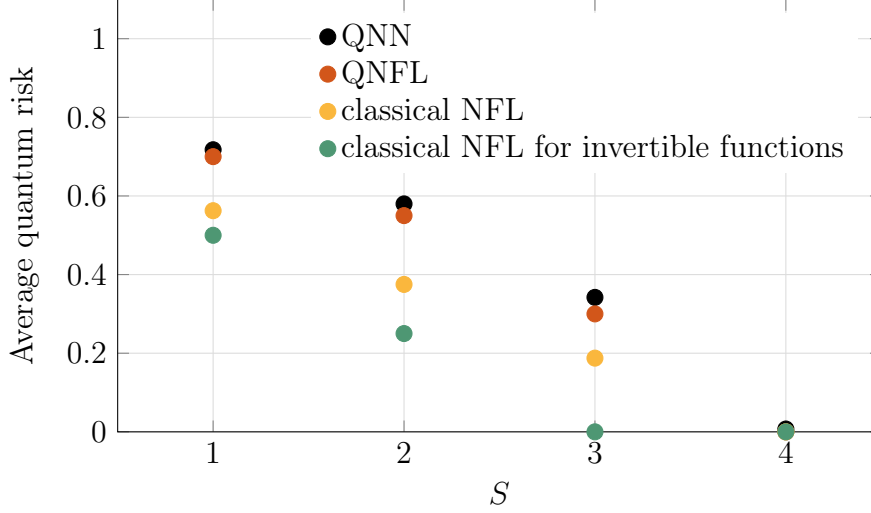


Figure 5.1.: **NFL theorem bounds.** The figure shows the average quantum risk based on the classical NFL theorem, the classical NFL theorem for invertible functions and the QNFL theorem for  $|X| = 4$  or respectively for a quantum data set of four state pairs. Further numerical results of the risk of a under the same conditions trained  $\otimes$  DQNN are plotted.

### Numerical results

At this point, three bounds on the training success of a learning algorithm are stated: the NFL theorem, the NFL theorem for invertible functions and the QNFL theorem. In the following, we not only want to compare these three results but also numerical results based on the DQNN algorithm introduced in Chapter 4.

In Figure 5.1, where the bounds are plotted, it becomes clear that the behaviour of NFL bound for invertible functions is very similar to the standard classical NFL theorem. The different slope shows the additionally available information when assuming the invertibility of the unknown function.

Further, one can see that the QNFL bound gives a stronger lower bound than its classical analogue. A paragraph in [100] describes the case of a single qubit as a good example to gain intuition: one could be wrong and assume that we can completely determine the action of  $Y$  when using just a single training pair  $\{|\phi_x^{\text{in}}\rangle, Y|\phi_x^{\text{in}}\rangle\}$ . Hence, the complement of the subspace  $\mathcal{K}$  spanned by  $|\phi_x^{\text{in}}\rangle$  has to be mapped to the complementary subspace determined by  $Y|\phi_x^{\text{in}}\rangle$ . However, this does not apply for the reason that there is nevertheless the freedom of a phase. One can get a feeling of the big impact of this freedom when noticing that we evaluate the risk averaged over the Hilbert space. In the here described case this is the average over all superposition inputs  $\alpha|\phi_x^{\text{in}}\rangle + \beta|\phi_x^{\text{in}}\rangle$ . Therefore the freedom of a phase affects the action of  $Y$  on nearly all inputs. In the classical case the analogue would be a function  $f$  on  $\{0, 1\}$ . Having the information of a single training pair, the action of  $f$  is determined for half of the inputs. Hence for half of

the cases, we are 100% sure. For the other input we can guess with being correct 50% of the time. In total, our hypothesis predicts the right output in 75% of the cases.

We so far only compared the QNFL theorem to the classical counterpart. As mentioned, the QNFL theorem is a bound on the training success of the in Chapter 4 described training algorithm. Consequently, Figure 5.1 also depicts the risk for an exemplary DQNN performed with the algorithm presented in Section 4.5. Therefore we choose a unitary  $Y$  uniformly at random, build 4 data pairs and train the QNN in  $r_T = 1000$  steps using  $S = 1, \dots, 4$  supervised training pairs. After the training,  $C = 10$  randomly chosen input states  $|\phi_x^R\rangle$  are used to evaluate the quantum risk by empirical average, namely

$$R_Y(U) = \frac{1}{C} \sum_{x=1}^C \|Y |\phi_x^R\rangle \langle \phi_x^R| Y^\dagger - U |\phi_x^R\rangle \langle \phi_x^R| U^\dagger\|_1^2,$$

where the unitary  $U$  describes the trained DQNN. The values describing the behaviour of the DQNN in Figure 5.1 are gained by averaging the risks  $R_Y(U)$  for 10 different unitaries  $Y$ .

We can note that the numerical results of the DQNN are close to achieving the QNFL bound. Since empirical averages are included in the process of evaluating the quantum risk and the DQNN is not trained to the maximum value of the training loss function, the slight discrepancy seen in Figure 5.1 was to be expected.

### Comment on orthonormal training pairs

We want to conclude the chapter with a short but essential note on using orthonormal training pairs, i.e.  $(|\phi_x^{\text{in}}\rangle, Y |\phi_x^{\text{in}}\rangle)$  with  $\langle \phi_x | \phi_k \rangle = \delta_{xk}$ , for learning an unknown unitary operation since this case can lead to problems. This circumstance can be explained with the above chosen quantum risk function, hence a remark at this point is suitable.

We want to start this discussion with the known case of not orthonormal  $|\phi_x^{\text{in}}\rangle$ . The aim is to train the network until it is  $U = Y$ , but indeed it is always possible that  $U$  is only equal to  $Y$  up to a global phase, namely  $U = Y \exp^{i\theta}$ . Applying the in Equation (5.2) defined risk function makes this clear, since

$$\sum_x \|Y |\phi_x^{\text{in}}\rangle \langle \phi_x^{\text{in}}| Y^\dagger - U |\phi_x^{\text{in}}\rangle \langle \phi_x^{\text{in}}| U^\dagger\|_1^2 = 0.$$

However, in the case of orthonormal training pairs, also local phases are possible, i.e. cases exist where the phase  $\theta_x$  depends on the training pair, namely  $U |\phi_x^{\text{in}}\rangle = \exp^{i\theta_x} Y |\phi_x^{\text{in}}\rangle$ . This by far worse case when learning a unitary operation is only possible, if the training pairs are orthogonal. To make this clear we can study the following case: the operation we want to learn is the identity,  $Y = \mathbb{1}$  and the training pairs are basis states  $\{|\phi_x^{\text{in}}\rangle\}_x = \{|0\dots 0\rangle, |0\dots 1\rangle, \dots, |1\dots 1\rangle\}$ . If global phases come across the risk function cannot tell the difference between the identity

## 5. No free lunch theorem

and the diagonal matrix with local phases written as  $U = \text{Diag}(\exp^{i\theta_1}, \dots, \exp^{i\theta_d})$ . We can reason from this short discussion that orthonormal training pairs should be avoided in training algorithms as described above.

To summarise, the QNFL theorem gives us a tool to review the in Chapter 4 presented DQNN algorithm. This algorithm uses training data for characterising a unitary operation. Here, the information for updating the DQNN parameters was based on training data pairs in the form of input states and desired output states. In the following chapter, we extend this supervised ansatz to the usage of the training data's possible underlying graph structure.



# 6

---

## Training with graph-structured quantum data

---

Graph-structured data is omnipresent throughout the social and natural sciences. In nearly every scientific area graphs, sets of in some way connected or related objects, are indispensable. These correlations, called edges, between the objects, named vertices, can be for example from spatial, temporal or causal nature. Whether it comes to representing computational devices, studying social media platforms [301], sampling road networks [302] or learning interactions between proteins [32], graphs provide a theoretical framework to describe these and many more complex systems in a helpful way [303].

Consequently, approaches including graph-structured data are widely spread in the area of classical *machine learning* (ML). At one hand unsupervised learning algorithms [304–307] exist. Such methods exhibit good learning behaviour with well-encoded graph structure and aim tasks like finding missing graph connections or labels. Further, there are semi-supervised methods [173, 308]. These use loss functions with one part making use of the graph structure of the data and another supervised loss term using supervised training pairs. Further, also methods encoding the graph structure directly in the representations, called graph convolutional neural networks [309–312], exist. A summary of some common classical approaches is presented in Section 6.2.

In the field of *quantum machine learning* (QML), research on how to work with graph-structured classical and quantum data was done as well. This includes quantum algorithms for classical graph-structured data like quantum walks [313], a quantum analogue of the random walk. Further, the usage of graph-structured data in quantum convolutional neural networks [270, 271] has received much attention.

Nevertheless, when it comes to using graph structure for *quantum neural networks* (QNNs), previous work has mainly focused on building the graph structure into

## 6. Training with graph-structured quantum data

the QNN [269, 270]. However, we present in this chapter an approach on how to make an arbitrary, not to the graph structure of the data aligned, QNN learn the graph structure of noisy and unreliable quantum data sources during training [99]. Based on the *dissipative quantum neural network* (DQNN) architecture presented in Chapter 4 we not only introduce a new loss function and suiting update rules for DQNNs, but also present examples that prove that the learning performance can be improved by learning and considering the graph structure in contrast to a simple supervised learning ansatz. We will see that exploiting this additional information can be highly beneficial, especially when data for supervised learning is rare but information on the structure of the problem is available.

To introduce the reader to the topic, we start this chapter with the basic graph theory definitions in Section 6.1 and give a survey of problems and methods in the field of classical ML on graphs Section 6.2. It follows a general discussion of quantum sources with graph structure, see Section 6.3, and the presentation of appropriate information-theoretic loss functions in Section 6.4 for their characterisation according to the approach of [99]. Further, we describe how to adjust the training algorithm described in Section 4.3 to utilise the new loss functions, see Section 6.7. We conclude this chapter with some numerical results: the classical simulation in Section 6.6, as well as the training on actual quantum computers Section 6.7, show the benefits of using graph information.

### 6.1 Basic definitions

This section will present the most fundamental definitions of graph theory. For a complete introduction to this field, we point to [314–316].




Figure 6.1.: **Exemplary graph.** This connected graph  $G_1$  consists of eight vertices  $V(G_1) = \{v_1, \dots, v_8\}$ , which are represented as heptagons and connected through the edges  $E(G_1) = \{e_{13}, \dots, e_{56}\}$ . The vertices  $v_2$  and  $v_4$  are labelled with labels  $\xi_2$  and  $\xi_4$ .

A *graph*  $G\{V, E\}$  is defined through the *vertex set*, also called *node set*,  $V(G) = \{v_1, \dots, v_N\}$ , and the *edge set*  $E(G) = \{e_{wx}\}_{w,x}$ . An edge  $e_{wx}$  connects vertices  $v_w$  and  $v_x$  if they are neighbours, referred to as  $v_w \sim v_x$ . These connections denote usually the relationship of the vertices, describes for example spacially, causal or temporal closeness of the vertices. The number of neighbours of a vertex  $v_x$  is called the *degree* of  $v_x$ .

## 6.2. Classical machine learning with graph-structured data

A *path* of  $k$  vertices is a sequence of  $k$  distinct vertices such that consecutive vertices are adjacent. We further denote a graph as *complete* if there is an edge between every pair of vertices. This definition should be not confused with a *connected* graph, where simple every pair of vertices can be connected with a path.

To be useful for ML graph-structure has to be encoded into continuous low-dimensional representations. A helpful tool is the *adjacency matrix*. For a graph with  $N$  vertices this  $N \times N$  matrix is defined by

$$A = \begin{cases} A_{wx} = 1 & \text{if there is an edge } e_{wx} \\ A_{wx} = 0 & \text{if there is no edge} \\ A_{xx} = 0. \end{cases}$$

The graph depicted in Figure 6.1 can be described with an adjacency matrix

$$A(G_1) = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

Note that in this work, we only discuss so-called *simple graphs*, where at most one edge between a pair of vertices is allowed. Further, we only study *undirected graphs*, where the existence of an edge  $e_{wx}$  is equivalent to the presence of an edge  $e_{xw}$ . Moreover, we assume that the relevance or meaning of all edges is equal. In contrast, there is the concept of *weighted graphs*, where the adjacency matrix' entries can be arbitrary values rather than just  $\in \{0, 1\}$ .

Every vertex can optionally be assigned to a *label*. Such a label can stand for a category, a number, a user profile in a social network, or some quantum data, to name a few examples. In the classical case, labels often are formulated as binary *label vectors*  $\xi(v_x)$ .

## 6.2 Classical machine learning with graph-structured data

Exploiting the graph structure of complex data sets has significant potential for new scientific breakthroughs. However, the challenge is to activate this potential by choosing the right techniques. Not only network analysis methods [317] but also ML approaches [175] were utilised for this.

ML techniques are often categorised into supervised tasks with the aim of getting the desired output given an input, or unsupervised tasks engaging with the assignment of learning patterns. Also semi-supervised tasks exist, which can be seen as a combination of both approaches. In the field of ML on graphs, these labels are used as well, just as mentioned in the introduction of this chapter. However, due to the characteristics of graphs, further (sub-)categories turned to

## 6. Training with graph-structured quantum data

be reasonable and valuable. We will present a selection of them shortly in the following paragraphs. This will supply the reader with a short but comprehensive overview of the field of ML on graphs. For a more detailed and comparing review we point to [175,318,319]. Note further that there exist methods, preparing graphs for such algorithms, for example by removing unnecessary neighbors [320].

### Node classification

A widespread problem category is *node classification* [318,321–323], where the aim is to use the information of the graph to find missing vertex labels. Imagine building a graph where the vertices refer to publications and the edges symbolise citations two papers have in common. If we then label some of the vertices with topics treated in the according to publication, it is possible to forecast missing labels by exploiting the graph structure [324]. Since usually the training set is built of the full graph (including the labelled and unlabelled vertices) and the existing labels, vertex classification is often classified as a semi-supervised task.

### Link prediction

Not only absent labels but also missing edges can be the focus of a problem. A well-known example for *link prediction* [318,325,326] can be found in polypharmacy: the usage of drug combinations is common for beating complex or co-existing diseases but also offers a higher risk of side effects. Data on these effects is usually scarce. However, [32] presents an approach named *Decagon* for modelling them. It uses a graph with protein-protein, drug-protein and drug-drug (polypharmacy) interactions, including many different edge types that stand for distinct side effects. The algorithm is evidentially able to find missing edges and therefore predict side effects of drug combinations. In the same way as vertex classification, relation prediction is often categorised as semi-supervised.

### Community detection

Whereas in the last two categories, missing graph data information was gained through the learning algorithm using some supervised data in the vertex labels, *community detection* [327] is an unsupervised process. The algorithms only input is the graph  $G = (V, E)$  itself. In many contexts, graphs naturally are clustered: vertices included in such a cluster are much more likely to be connected with other included vertices than with those from outside the cluster. An example is the connection of user profiles on social media platforms: these will be likely clustered, for instance, by the users home location. Algorithms finding such clusters in graphs can be used, for example, to discover fraudulent anomalies in financial transaction networks given the histories of transactions of people using the networks [328].

### Graph classification

On the contrary to the three above mentioned categories *graph classification* [329] is not based on one single graph but is instead feed with different graphs in order to categorise them. For example, we assume a given graph-based representation of the structure of some molecules. The algorithm can be trained such that it reads the graph and classifies the molecule in some way, for instance, by its toxicity [330].

### Graph-convolution networks

The last method we want to mention in this short overview is directly encoding the graph structure in the representations. These so-called *graph-convolution networks* (GCNs) [309–312] are one of the most prominent graph deep learning models and especially useful when other graph-based ML methods face problems. The techniques can also be used on data types that are not structured initially as a graph, such as image or text data, and graph-structured data with very complex patterns. A concrete example is building a representation of video input as a space-time region graph and recognising actions in these videos, for example like "opening a book" [331]. A review of GCNs can be found in [332].

## 6.3 Quantum graph-structured data

In the last section, we gained an insight into the success of classical graph-structured data processed in classical ML algorithms. Nevertheless, not only these achievements are the motivation to study the use of quantum graph-structured data. The properties of the data itself prompt this kind of ansatz: quantum data, produced by structured devices, will always be structured since spatial and causal arrangements lead to different correlations between the states. In the following, we give some examples of how graph-structured data can emerge. We close this section by developing a notation describing the graph structure of the quantum data set in a way that is useful for the training of QNNs.

### Motivational examples

One can imagine a set of  $N$  quantum information processors represented by the vertices of a graph. We assume we know in which way these processors are arranged. For every processor, some of the other processors are closer than the rest. We can relate particularly close processors with edges between the concerning vertices encoding the correlations between the processors. We further assume we not only know the structure of the processors, but also have access to a training data set of  $S < N$  outputs and inputs  $\{(\rho_x^{\text{in}}, \rho_x^{\text{out}}) \mid x = 1, 2, \dots, S\}$  defined as follows: we know that processing the state  $\rho_x^{\text{in}}$  with the processor  $x$  can lead to the state  $\rho_x^{\text{out}}$ . In other words, we own a training pair consisting of an input and output state for some of the processors, but not all of them. The aim would be now to learn the input and output relations for all the  $N$  processors. This scenario can be seen as

## 6. Training with graph-structured quantum data

the foundation of various physically relevant situations. The quantum NISQ device clusters [53] are one example where the quantum data is in some way spacially graph-structured.

An other example is structure due to time steps. Imagine simulating a quantum system with Hamiltonian  $H \in \mathcal{B}(\mathcal{H})$  for some period of time in steps  $t \in \{0, \epsilon, 2\epsilon, \dots, (N-1)\epsilon\}$ . Based on an initial state  $|\rho^{\text{in}}\rangle \in \mathcal{H}$  the states  $|\psi_t\rangle \equiv e^{itH/\hbar} |\rho^{\text{in}}\rangle$  evolve. Such a structure can also be interpreted as a line graph with  $N$  vertices.

### Notation

After motivating graph-structured quantum data, we introduce some convenient notation to describe such in the following. For this, we tie in with Section 4.3 and assume having access to one or more quantum devices producing uncharacterised quantum states. We further suppose that the quantum information device can be described by a completely positive map  $\mathcal{E}$  mapping from an input state  $\rho_x$  to an output state  $\sigma_x = \mathcal{E}(\rho_x)$ . The device can be described with  $\{(p_x, \rho_x)\}_{x \in V}$ , where  $\rho_x$  occurs with probability  $p_x$ .

Additionally to the setting in Chapter 4 we presume the quantum states  $\rho_x$  are linked to the vertices of a graph  $G = (V, E)$ , namely

$$\rho : V \rightarrow \mathcal{D}(\mathcal{H}).$$

where  $V$  denote the vertices of the graph and  $\mathcal{D}(\mathcal{H})$  the density matrices on  $\mathcal{H}$ . More specifically, we assume that the graph structure is associated with the device's outputs.

We further exhibit the graph-structure with an edge set  $E$ . In this sense two states  $\rho_w$  and  $\rho_x$  which are *information-theoretic close*, defined by some metric  $d(\rho_w, \rho_x) \sim \epsilon$ , are connected with an corresponding edge  $(w, x) \in E$ . We discuss the choice of the metric in Section 6.4. These edges can be described in an adjacency matrix  $A$ .

Additionally to this  $N \times N$  matrix we have a training data set on hand. Therefore this setting can be categorised as semi-supervised. The training data set is structured as

$$\left\{ (\rho_1, |\phi_1^{\text{SV}}\rangle \langle \phi_1^{\text{SV}}|), \dots, (\rho_S, |\phi_S^{\text{SV}}\rangle \langle \phi_S^{\text{SV}}|), \rho_{S+1}, \dots, \rho_N \right\}, \quad (6.1)$$

where the first  $S$  entries are the quantum state pairs describing the supervised data (as in the supervised setting discussed in Chapter 4) and the remaining  $N - S$  entries describe unsupervised vertices without information about the output. Throughout this chapter, we will call the vertices labelled with the input state and the desired output state as *supervised* and the remaining vertices as *unsupervised*. An example containing one supervised vertex and two unsupervised vertices is depicted in Figure 6.2.

We aim to find the output states of the unsupervised vertices by exploiting the

supervised vertices and the graph structure. Hence we can categorise the problem as a semi-supervised node classification question.

For the validation process of our examples presented in Section 6.6 we assume moreover to have access to the *overall* data set, i.e.

$$\left\{ (\rho_1, |\phi_1^{\text{SV}}\rangle \langle \phi_1^{\text{SV}}|), \dots, (\rho_S, |\phi_S^{\text{SV}}\rangle \langle \phi_S^{\text{SV}}|), \right. \\ \left. (\rho_{S+1}, |\phi_{S+1}^{\text{USV}}\rangle \langle \phi_{S+1}^{\text{USV}}|), \dots, (\rho_N, |\phi_N^{\text{USV}}\rangle \langle \phi_N^{\text{USV}}|) \right\}.$$

Using this additional data, we can test if the new graph loss function, presented in the next section, increases the generalisation behaviour of the DQNN training algorithm. In the use case, this additional data is not necessary, and the data in the form of Equation (6.1) is sufficient. Note that we assume that the training and validation data consist of pure states in the following, but mixed states are possible in general as well.




Figure 6.2.: **Graph and QNN.** A line graph with  $N = 3$  vertices whereof  $S = 1$  is supervised is depicted on the left-side of the figure. The labels of the vertices are either input states (unsupervised vertex) or pairs of input and output states (supervised vertex). The dimension of the input state should match the input layer of the QNN, the dimension of the supervised states the output layer, respectively.

## 6.4 Loss functions

As in Chapter 4 we use pairs of quantum states as training data to learn unknown quantum information processes. In addition to that, we include here information about the problem's graph structure. The task is now to find a way that the in Chapter 4 presented DQNN can learn and generalise from this graph data. Hence in the following, we will guide through different loss functions used in the proposal of [99]. To remind the reader, we describe the learning architecture we train with the function  $\mathcal{E}$ , i.e.  $\rho^{\text{out}} = \mathcal{E}(\rho^{\text{in}})$ .

## 6. Training with graph-structured quantum data

### Training loss

We already discussed how to train a DQNN with pure supervised states using the fidelity in Section 4.3, namely with the *supervised loss*

$$\mathcal{L}_{\text{SV}} \equiv \frac{1}{S} \sum_{x=1}^S \langle \phi_x^{\text{SV}} | \mathcal{E}(\rho_x^{\text{in}}) | \phi_x^{\text{SV}} \rangle.$$

In addition, we define here a new loss in order to exploit the graph structure of the output states of the device. These output states are, in general, mixed. Even though the fidelity is defined for mixed states, the computation complexity required to evaluate it is immense. Therefore, we choose a different measure here, namely the *Hilbert-Schmidt distance*

$$d_{\text{HS}}(\rho, \sigma) \equiv \text{tr}((\rho - \sigma)^2).$$

Keep in mind that the Hilbert-Schmidt distance reaches its minimum if the two compared states resemble. In contrast, the fidelity reaches the maximum, the value 1, in this case. See Section 3.1 for a more detailed discussion of the fidelity and the Hilbert-Schmidt distance.

We use the adjacency matrix  $A$  of the graph  $G$  to include the graph structure information during the learning process and define the *graph-based loss* as

$$\mathcal{L}_G \equiv \sum_{w,x \in V} [A]_{wx} d_{\text{HS}}(\mathcal{E}(\rho_w^{\text{in}}), \mathcal{E}(\rho_x^{\text{in}})),$$

where  $[A]_{wx}$  denotes the matrix element of  $A$  relating to vertices  $v$  and  $w$ .

The full *training loss function* is now specified as the combination of supervised and graph-based loss, with the graph part controlled by a Lagrange multiplier  $\gamma$ :

$$\mathcal{L}_{\text{SV+G}} = \mathcal{L}_{\text{SV}} + \gamma \mathcal{L}_G.$$

Defining the training loss in that way forces the network to map supervised input states to the desired output states and also considers the graph structure of the data.

In the same way as represented in Section 4.3 where we described the training of a neural network without using graph-structure, the training task is to maximise the training loss function, here  $\mathcal{L}_{\text{SV+G}}$  with  $\gamma \leq 0$ . The maximum depends on the Lagrange multiplier, and by tuning it, we can decide the importance of the graph structure while training.

### Validation loss

We aim the quantum neural network to lead to the right output state for a given input state in the same way as in Chapter 4 and therefore use the equivalent



validation loss, namely

$$\mathcal{L}_{\text{USV}} = \frac{1}{N-S} \sum_{x=S+1}^N \langle \phi_x^{\text{USV}} | \mathcal{E}(\rho_x^{\text{in}}) | \phi_x^{\text{USV}} \rangle,$$

to check the trained quantum neural networks behaviour using the unsupervised data pairs. We will use this function to study numerical examples in Section 6.6.

## 6.5 Training algorithm

For training the DQNN, including the graph information, we use an algorithm of the same structure as described in Section 4.3. The only difference is the new training loss function  $\mathcal{L}_{\text{SV+G}} = \mathcal{L}_{\text{SV}} + \gamma \mathcal{L}_{\text{G}}$  leads to new update rules which we formulate in Hermitian matrices  $K_{j,\text{SV+G}}^l(s)$  for  $j$ th qubit in  $l$ th layer. Following the discussions in Chapter 4 we recognise that all calculations are linear in the loss function, and therefore we can make the ansatz

$$K_{j,\text{SV+G}}^l(s) = K_{j,\text{SV}}^l(s) + \gamma \cdot K_{j,\text{G}}^l(s),$$

where  $K_{j,\text{SV}}^l(s)$  denotes the update matrix derived in Section 4.3. It remains to derive the unsupervised update matrix  $K_{j,\text{G}}^l(s)$  to clarify the training algorithm when using graph-structured quantum data for training a DQNN.

**Proposition 6.1.** *The update matrix for training a QNN with a graph structure between the output states  $\{\rho_w^{\text{out}}, \rho_x^{\text{out}}\}$  encoded via the adjacency matrix  $[A]_{wx}$  (and without any supervised states) is*

$$K_{j,\text{G}}^l(s) = 2^{m_{l-1}+1} i\eta \sum_{v \sim w} [A]_{wx} \text{tr}_{\text{rest}}(M_{j\{w,x\}}^l(s)),$$

where

$$\begin{aligned} M_{j\{w,x\}}^l(s) = & [U_j^l(s) U_{j-1}^l(s) \dots U_1^l(s) ((\rho_w^{\text{in}} - \rho_x^{\text{in}}) \otimes |0 \dots 0\rangle_1 \langle 0 \dots 0|) \\ & U_1^{l\dagger}(s) \dots U_{j-1}^{l\dagger}(s) U_j^{l\dagger}(s), \\ & U_{j+1}^{l\dagger}(s) \dots U_{m_{\text{out}}}^{\text{out}\dagger}(s) (\mathbb{1}_{\text{in,hidden}} \otimes (\rho_w^{\text{out}} - \rho_x^{\text{out}})) U_{m_{\text{out}}}^{\text{out}}(s) \dots U_{j+1}^l(s)]. \end{aligned}$$

*Proof.* The derivation is analogous to the proof of Proposition 4.1. Given the graph-based loss

$$\mathcal{L}_{\text{G}} \equiv \sum_{w,x \in V} [A]_{wx} d_{\text{HS}}(\mathcal{E}(\rho_w^{\text{in}}), \mathcal{E}(\rho_x^{\text{in}})),$$

we take the derivative with respect to the step parameter. We get

$$\frac{d\mathcal{L}_{\text{G}}(s)}{ds} = \lim_{\epsilon \rightarrow 0} \frac{\mathcal{L}_{\text{G}}(s + \epsilon) - \mathcal{L}_{\text{G}}(s)}{\epsilon}$$

## 6. Training with graph-structured quantum data

$$\begin{aligned}
&= 2i \sum_{w,x \in V} [A]_{wx} \operatorname{tr} \left( (\mathbb{1}_{\text{in, hidden}} \otimes (\rho_w^{\text{out}} - \rho_x^{\text{out}})) \right. \\
&\quad \left( [K_{m_{\text{out}}}^{\text{out}}, U_{m_{\text{out}}}^{\text{out}} \dots U_1^1] U_1^{1\dagger} \dots U_{m_{\text{out}}}^{\text{out}\dagger} \right) + \dots \\
&\quad \left. + U_{m_{\text{out}}}^{\text{out}} \dots U_2^1 [K_1^1, U_1^1 (\rho_w^{\text{in}} - \rho_x^{\text{in}}) U_1^{1\dagger}] U_2^{1\dagger} \dots U_{m_{\text{out}}}^{\text{out}\dagger} \right) \\
&= 2i \sum_{w,x \in V} [A]_{wx} \operatorname{tr} \left( [U_{m_{\text{out}}}^{\text{out}} \dots U_1^1 (\rho_w^{\text{in}} - \rho_x^{\text{in}}) U_1^{1\dagger} \dots U_{m_{\text{out}}}^{\text{out}\dagger}, \right. \\
&\quad \left. (\mathbb{1}_{\text{in, hidden}} \otimes (\rho_w^{\text{out}} - \rho_x^{\text{out}})) \right] K_{m_{\text{out}}}^{\text{out}} + \dots + [U_1^1 (\rho_w^{\text{in}} - \rho_x^{\text{in}}) U_1^{1\dagger}, \\
&\quad U_2^{1\dagger} \dots U_{m_{\text{out}}}^{\text{out}\dagger} (\mathbb{1}_{\text{in, hidden}} \otimes (\rho_w^{\text{out}} - \rho_x^{\text{out}})) U_{m_{\text{out}}}^{\text{out}} \dots U_2^1] K_1^1 \\
&= 2i \sum_{w,x \in V} [A]_{wx} \operatorname{tr} \left( M_{m_{\text{out}}\{w,x\}}^{\text{out}}(s) K_{m_{\text{out}}}^{\text{out}}(s) + \dots M_{1\{w,x\}}^1(s) K_1^1(s) \right),
\end{aligned}$$

where

$$\begin{aligned}
M_{j\{w,x\}}^l(s) &= [U_j^l(s) U_{j-1}^l(s) \dots U_1^1(s) ((\rho_w^{\text{in}} - \rho_x^{\text{in}}) \\
&\quad \otimes |0 \dots 0\rangle_1 \langle 0 \dots 0|) U_1^{1\dagger}(s) \dots U_{j-1}^{l\dagger}(s) U_j^{l\dagger}(s), \\
&\quad U_{j+1}^{l\dagger}(s) \dots U_{m_{\text{out}}}^{\text{out}\dagger}(s) (\mathbb{1}_{\text{in, hidden}} \otimes (\rho_w^{\text{out}}(s) - \rho_x^{\text{out}}(s))) U_{m_{\text{out}}}^{\text{out}}(s) \dots U_{j+1}^l(s)].
\end{aligned}$$

As a next step we expand  $K_{j,G}^l(s)$  using Pauli matrices

$$K_{j,G}^l(s) = \sum_{\alpha_1, \dots, \alpha_{m_l-1}, \beta} K_{j, \alpha_1, \dots, \alpha_{m_l-1}, \beta}^l(s) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^{\alpha_{m_l-1}} \otimes \sigma^\beta).$$

Since  $d_s \mathcal{L}_G(s)$  is linear in the coefficients  $K_{j, \alpha_1, \dots, \alpha_{m_l-1}, \beta}^l(s)$  we can solve

$$\min_{K_{j, \alpha_1, \dots, \beta}^l} \left( \frac{d\mathcal{L}_G(s)}{ds} - \lambda \sum_{\alpha_1, \dots, \beta} K_{j, \alpha_1, \dots, \beta}^l(s)^2 \right),$$

analogously to the proof in Proposition 4.1 using the Lagrange multiplier  $\lambda \in \mathbb{R}$ . Inserting the resulting coefficients

$$K_{j, \alpha_1, \dots, \beta}^l = \frac{i}{\lambda} \sum_{w,x \in V} [A]_{wx} \operatorname{tr}_{\alpha_1, \dots, \alpha_{m_l-1}, \beta} \left( \operatorname{tr}_{\text{rest}} (M_{j\{w,x\}}^l) (\sigma^{\mu_1} \otimes \dots \otimes \sigma^\eta) \right)$$

gives the desired expression for the update matrices, namely

$$K_{j,G}^l(s) = 2^{m_l-1+1} i \eta \sum_{w,x \in V} [A]_{wx} \operatorname{tr}_{\text{rest}} (M_{j\{w,x\}}^l(s)),$$

where  $\eta = 1/\lambda$  is the learning rate. □

Using the above-reasoned ansatz

$$K_{j,\text{SV+G}}^l(s) = K_{j,\text{SV}}^l(s) + \gamma \cdot K_{j,\text{G}}^l(s)$$

we end up with the update matrix

$$K_{j,\text{SV+G}}^l(s) = \frac{2^{m_{l-1}} \eta_i}{S} \sum_{x=1}^S \text{tr}_{\text{rest}} (M_{j\{x\}}^l(s)) + \gamma 2^{m_{l-1}+1} i \eta \sum_{w,x \in V} [A]_{wx} \text{tr}_{\text{rest}} (M_{j\{w,x\}}^l(s)),$$

for maximising the training loss  $\mathcal{L}_{\text{SV+G}} = \mathcal{L}_{\text{SV}} + \gamma \mathcal{L}_{\text{G}}$ , where  $\gamma < 0$  is the parameter to tune the influence of the graph structure.

## 6.6 Classical simulation

In the last section, we derived the update rule for the training loss, including the graph's information. Using these results and the algorithm described in Section 4.3 we can simulate the training on a quantum computer using QuTip [275], a quantum toolbox in Python. The results are presented analogously to the training of the DQNN without graph structure in Section 4.5. The code can be found at [276].

As already mentioned in Section 4.5 we are limited to small quantum systems due to the exponential scaling of the Hilbert space dimension with number of qubits. Despite this constraint, we will present three numerical studies following [99].

### Example I: connected clusters


The first graph we study is a connected graph of degree  $N = 8$ . The vertices have pairs of quantum states, including an input and output state, as labels. The first four vertices  $v_1 \dots v_4$  form a cluster, as well as the vertices  $v_5 \dots v_7$ . Vertex  $v_8$  connects the two clusters, see Figure 6.3a. The output states are chosen in a way that connected vertices are associated to *closer* output states in the sense of the Hilbert-Schmidt distance. The input states  $|\phi_x^{\text{in}}\rangle$  are random 3-qubit states built via a normal (Gaussian) distribution and are not linked to the graph structure in any way. Note that the coefficients depicted in Figure 6.3a are only recorded to three decimal places.

The graph structure, saved in an adjacency matrix, is the foundation of the training loss  $\mathcal{L}_{\text{G}}$ . As introduced above, we also include a supervised learning part, packed in  $\mathcal{L}_{\text{SV}}$ . For this purpose, we assume that  $S$  of the state pairs are used for training. In Figure 6.3a these  $S = 3$  supervised vertices are shaded.


Figure 6.3b depicts the validation loss during  $r_T = 1000$  training epochs in the supervise case ( $\gamma = 0$ ) supervised and semi-supervised (i.e. supervised plus graph-based) case ( $\gamma = -0.5$ ) with  $S = 3$ . Although we do not reach a fidelity of 1, the plot clarifies that the learning algorithm performs better when the graph structure is exploited.

We already observed that interpolation of the action on the supervised vertices on the unsupervised vertices is possible after training the network with three of


## 6. Training with graph-structured quantum data



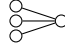
(a) Graph with labels.




(b) Validation loss during training with  $S = 3$  supervised training pairs.




(c) Validation loss after training with different  $S$  averaged over 30 sets.

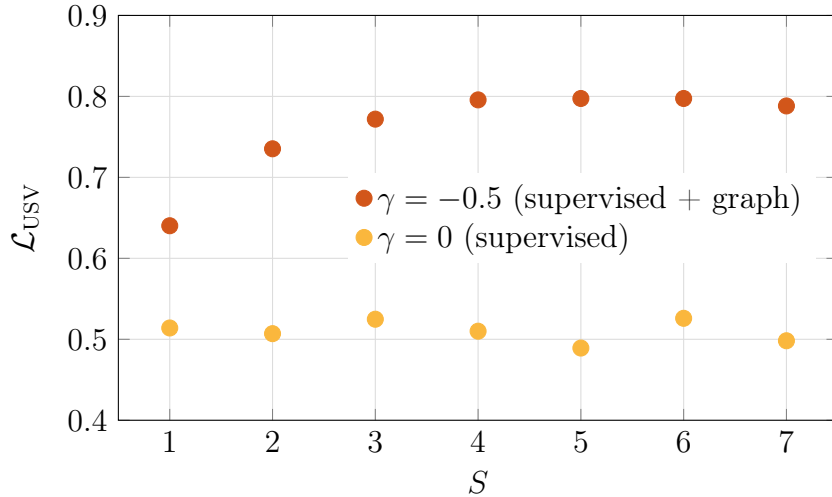
Figure 6.3.: **Connected clusters.** This figure compares the training (b) and generalisation behaviour (c) of optimising a  DQNN ( $r_T = 1000$  epochs,  $\epsilon = 0.01$ ) with and without using the graph structure represented in (a).



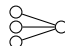
(a) Graph with labels.



(b) Validation loss during training with  $S = 3$  supervised training pairs.



(c) Validation loss after training with different  $S$  averaged over 30 sets.

Figure 6.4.: **Line.** This illustration draws the comparison of the training (b) and generalisation behaviour (c) of optimising a  DQNN ( $r_T = 1000$  epochs,  $\epsilon = 0.01$ ) trained with and without using the graph structure (a).

## 6. Training with graph-structured quantum data

the eight data pairs. We test how the number of supervised vertices  $S$  affects the training process for a generalisation study. Therefore we randomly chose  $S < N$  of the  $N = 8$  training pairs, trained the network in  $r_T = 1000$  steps and average the last values of the loss functions over 10 completely independent training attempts. The results are depicted in Figure 6.3c and show that for  $S \geq 2$  the graph-based loss term optimises the training significantly.

### Example II: Line

In the second example,  $N = 10$  pairs of quantum states are the labels of vertices aligned in a line graph, see Figure 6.4a: the states were chosen to be – according to the fidelity – evenly spaced along a line between the endpoints associated with  $|0\rangle$  and  $|1\rangle$ . In the same way, as in the connected-clusters data set, the input states  $|\phi_x^{\text{in}}\rangle$  are random 3-qubit states and the output states suit the graph structure. Again we assume three vertices to be supervised.

In Figure 6.4b the validation loss is depicted. The training loss involving the graph structure ( $\gamma = 0$ ) yields a much higher validation loss function during the training. As in the example above the graph, structure supports the training. The maximum value of the fidelity is not reached, but after only 376 training epochs, we exceed  $\mathcal{L}_{\text{USV}}(s) = 0.8$ .

In the generalisation study, see Figure 6.4c, it becomes clear that a validation loss of over 0.6 with only 5 of 10 supervised vertices can be achieved when the graph structure is exploited.

### Example III: Classic deep walk

Using the quantum states forming connected clusters or a line graph we could, as described above, observe that including the graph structure of the output states in the training process can increase the training success. However, the input states, fed into the QNN, were random states. In the following we describe an ansatz, where both, the input and the desired output states, are related to a graph.

We use a synthetic graph, motivated by the social distance attachment model in [333]. In contrast to the graphs used in this chapter before to every vertex not only labels are assigned, but also an embedding vector. The graph is built classically by randomly assigning labels to vertices such that the average number of labels per vertex is a constant. The number of possible labels is 8, this means we can describe the labels of every vertex  $x$  as a string of 8 binary numbers. Our graph contains 32 vertices. The average number of labels per vertex is 3.

To create edges between the labelled vertices, let  $d(\xi_w, \xi_x)$  denote the hamming distance between the label vectors  $\xi_w$  and  $\xi_x$  for vertices  $w$  and  $x$ . We denote with  $h$  the level of homophily and  $b$  the characteristic distance. In graph theory, homophily describes how often linked vertices have the same labels or similar features [334]. Further we generate an edge between two vertices  $w$  and  $x$  with


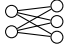


Figure 6.5.: **Deep walk.** The plot describes the generalisation behaviour of a  DQNN ( $r_T = 2000$  epochs,  $\epsilon = 0.01$ ) trained with and without using the graph structure of a graph with 32 vertices produced by a classical deep walk. Each data point demonstrates an averaged over 5 independent training attempts.

the probability

$$p_{wx} = \frac{1}{1 + [b^{-1}d(\xi_w, \xi_x)]^h}.$$

As higher we choose the value of  $h$ , the chance of vertices with similar labels are more likely connected.

As a last step we assign every vertex to an embedding vector  $\vec{e} = \{\alpha, \beta, \gamma, \delta\}$ . The embedding vector is computed using the method *DeepWalk* [304] using walks of length 1 and the number of walks per vertex as 10. These embedding vectors will be used to construct input quantum states.

We now construct the corresponding quantum data, namely the quantum input and output states of the vertices. We adopt the graph structure and change the labels in the following way. Every vertex is assigned to a pair of quantum states: a 2-qubit state  $|\phi_x^E\rangle$  based on the embeddings and a 3-qubit state  $|\phi_x^L\rangle$  build based on the labels.  $|\phi_x^E\rangle$  is the normalised version of the superposition  $\alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$ . To build  $|\phi_x^L\rangle$  we link the states  $\{|000\rangle, \dots, |111\rangle\}$  to the 8 possible labels. The output state assigned to a specific vertex is now built as the superposition of these basis states assigned to the vertex' labels. If for example a vertex has the labels 2, 3 and 8 the output state would be a superposition of the states  $|001\rangle, |010\rangle$  and  $|111\rangle$ .

We use  $|\phi_x^E\rangle$  as input states and  $|\phi_x^L\rangle$  as desired output states to train DQNNs. The generalisation analysis in Figure 6.5, shows that with  $3 \leq S \leq 10$  supervised vertices, the validation loss is lower when ignoring the problem's graph structure. For larger values of  $S$  both learning strategies are about equally good, which

## 6. Training with graph-structured quantum data

is discussed in the appendix, see Figure B.2 using a smaller amount of training rounds, since a study of the loss functions for all  $S$  and  $r_T = 2000$  exceeds the computational power of the authors of this thesis. The evaluation of the loss functions is plotted in the appendix as well, see Figure B.2.

### 6.7 NISQ device implementation

In the last section, we discussed training a DQNN using the newly introduced training loss function  $\mathcal{L}_{SV+G}$ . The results were done using the classical simulation presented in Section 4.5. We will now test the behaviour of this new loss function using the NISQ device implementation, denoted with  $\text{DQNN}_{\text{NISQ}}$  and discussed in Section 4.6.

As in Section 4.6 explained, the update of the parameter vector  $\omega_s$  is done in classical manner using the gradient of the training loss, i.e.  $\omega_{s+1} = \omega_s + d\omega$ , where  $d\omega = \eta \nabla \mathcal{L}_{SV+G}(\omega_t)$  and

$$\nabla_k \mathcal{L}_{SV+G}(\omega_s) = \frac{\mathcal{L}_{SV+G}(\omega_s + \epsilon e_k) - \mathcal{L}_{SV+G}(\omega_s - \epsilon e_k)}{2\epsilon} + \mathcal{O}(\epsilon^2).$$


$\mathcal{L}_{SV+G}$  function includes both the fidelity and the Hilbert-Schmidt distance. Therefore, additionally to the implementation presented in Section 4.6, we have to evaluate the Hilbert-Schmidt distance. This can be done by three evaluations of the *SWAP* test, since  $d_{\text{HS}}(\rho, \sigma) = \text{tr}(\rho^2) - 2\text{tr}(\rho\sigma) + \text{tr}(\sigma^2)$  and  $\text{tr}(\rho\sigma) = p \cdot c$  can be measured with the *SWAP* test [205], explained in Section 4.6.

The training of the  $\text{DQNN}_{\text{NISQ}}$  using the graph-structured quantum data is presented in Figure 6.6a (connected clusters) and Figure 6.6a (line). Despite the noise levels, we can observe that using the graph structure of the problem increases the validation loss reached after about 500 epochs and seems to lead to a more stable training.


The plots in Figure 6.6 are based on numerics presented in the work of [335]. We point the reader to this source for more experiments with graph-structured quantum data and NISQ devices and a comparison of the behaviour of  $\text{DQNN}_{\text{NISQ}}$  and QAOA with the semi-supervised loss function  $\mathcal{L}_{SV+G}$ .

In this chapter we presented a variation of the DQNN training algorithm presented in Chapter 4 by including knowledge of the graph structure of the training data into the learning process. We could see that this extension can increase the reached validation loss. In the next chapter, we will extend the DQNN ansatz in another direction and follow a generative adversarial approach where two DQNNs, a generator and a discriminator model, are trained in a competitive manner.





(a) Training with connected clusters data set.



(b) Training with line data set.

Figure 6.6.: **Training a  $\text{DQNN}_{\text{NISQ}}$  using graph-structured quantum data.** The figures depict the validation loss during training a  $\text{DQNN}_{\text{NISQ}}$  in 500 epochs with the datasets connected clusters (a), see Figure 6.3a, and line (b), see Figure 6.4a. In both training attempts  $S = 3$  training pairs are supervised.



# 7

---

## Quantum generative adversarial networks

---

In the preceding chapters, we worked with training data in the form of pairs of input and output quantum states in order to learn the relation between them. We defined loss functions that compared the output of a *quantum neural network* (QNN) with the desired output and demonstrated that is possible to update the network based on this and achieve good training results. The whole training algorithm rests upon the fact that we have access to related data pairs. In contrast, in the following we undertake the task of extending a set of quantum states to states which have similar properties.

We present a QNN training algorithm [102] which can not only learn the common characteristics of such data sets but also generates states with suiting requirements. The proposal is based on the *dissipative quantum neural networks* (DQNNs) presented Chapter 4. More precisely, we will use two DQNNs acting in competition, whereof one acts in a generative and the other in a discriminative manner, to learn from unlabelled quantum data.

In classical *machine learning* (ML) we generally distinguish between *discriminative* and *generative* models. The in Chapter 2 discussed example of classifying handwritten digits based on the MNIST dataset [148] is a typical discriminative task: the *neural network* (NN) learns in a supervised way to decide between the ten options. Such problems have in common that it is easy to find a proper training loss function which is aimed to be optimised during training [3]. In contrast, generative models produce data. Speaking in the example of handwritten digits, we would train a generative model to produce different “handwritten” digits from a random input. The downside of generative models is that these, due to the difficulty of approximating many probabilistic computations, are much more complicated to train than discriminative models.

## 7. Quantum generative adversarial networks

This changed with the arrival of *generative adversarial networks* (GANs) [147], consisting of a generative and discriminative part. The generative model generates samples by passing random noise through, for example, a multi-layer perceptron. The discriminative part, which can also be realised by a multi-layer perceptron, is trained to distinguish the data produced by the generator from the training data. On the other hand, the generative model aims to “fool” the discriminator. In this opposing way, it is possible to train the generator [147, 162, 336] and GANs have since found a lot of applications [336], ranging from classification or regression tasks [336–338] to the generation [339] and improvement [340] of images.

In the field of quantum machine learning (QML), such generative adversarial processes, referred to as *quantum generative adversarial networks* (QGANs) [341–348], were studied, as well. Some of these models are defined as a quantum-classical hybrid, include a quantum generator and a classical discriminator, and are used to learn from classical input data [346, 347]. Moreover, the authors of [342] present the usage of quantum or classical data and two quantum processors in an adversarial learning setting. Furthermore, the work of [341] designs quantum circuits for the generator and discriminator NNs and trains them using data which has two possible labels assigned. A further recent proposal is the *Entangling* QGAN [348]: here, the discriminative model has access to a state produced by the generator a training data set and a parametrised *SWAP* test is used to “entangle” the states.

Since we could observe excellent learning behaviour using the in Chapter 4 presented DQNNs in a supervised ansatz, in this chapter, we introduce a QGAN based on these types of QNNs. This so called *discriminative quantum generative adversarial network* (DQGAN) [102] is, in contrast to some of the above named proposals, a fully quantum architecture and is trained with quantum data. Different to the approach in [341] in the DQGAN algorithm only the discriminator has access to training data, and we only work with unlabelled quantum states as training data. In fact, the generative model gets a random quantum state as input, whereas the discriminator receives either the output of the generative model or a training data state as input. Since quantum data is always rare, this model aims to produce states with similar features as the training data. In that way, we generated more quantum data of the same kind, which could be, for example, useful to train other QNNs.

We start this chapter with a discussion of classical GAN methods in Section 7.1 since the quantum analogues are based on these. It follows a presentation of the already advertised DQGAN architecture, see Section 7.2. We discuss the data sets and loss functions used for training and validation of DQGAN in Section 7.3. Based on this, we formulate the update rules for training the DQGAN training algorithm in Section 7.4. We conclude with a classical simulation of the algorithm and show results of an implementation suitable for a NISQ device in Section 7.5 and Section 7.6.

## 7.1 Classical adversarial networks

In this section, we summarise the work of [147], who propose the first GAN model and explore the particular case where the generative model generates samples by passing random noise through a multi-layer perceptron. The discriminative model is built as a multi-layer perceptron as well. It has access to generator's synthetic samples and samples from the stack of training data. However, the generator part has no direct admission to the training data. In general, GANs are usually realised with multi-layer NNs consisting of convolutional or fully connected layers [336].

### Network architecture

For describing the proposal of [147] we call the generative model, capturing the data distribution and producing data,  $G$ . The other model, the discriminator, is denoted by  $D$  and estimates the probability  $D(x)$  that a given sample  $x$  is based on the training data set rather than produced by the generator  $G$ .

Here  $G$  takes a random input  $z$  with distribution  $p_{\text{in}}$  and outputs  $x$  with probability distribution  $p_G$ . On the other hand the probability distribution  $p_{\text{data}}$  describes sampling  $x$  from the training data. The second network  $D$  gets  $x$  either based on  $p_G$  or  $p_{\text{data}}$  as an input. These relations are depicted in Figure 7.1 for the case  $G(x) : [0, 1] \rightarrow [0, 1] \times [0, 1]$ ,  $D(x) : [0, 1] \times [0, 1] \rightarrow [0, 1]$ , and  $D(G(x)) : [0, 1] \rightarrow [0, 1]$ .




Figure 7.1.: **GAN**. The discriminative model  $D$  gets the input  $x$  either of the generator  $G$  or from the training data set.

### Training

For the training we assume that both models,  $G(x, \theta_G)$  and  $D(x, \theta_D)$  are implemented by multilayer perceptrons and parametrised by  $\theta_G$  and  $\theta_D$ , respectively. Both models are trained simultaneously but with a different aim. The training goal of model  $G$  is to maximise the probability of  $D$  making a mistake. The aim of  $D$  is to always make the correct distinction. Therefore the whole problem corresponds to a *minimax problem*: We train  $D$  to maximise the probability of assigning the correct label to both training examples and samples from  $G$ . Simultaneously we

## 7. Quantum generative adversarial networks

train  $G$  to minimise  $\log(1 - D(G(z)))$ , i.e. we train the generator to produce fake data which cannot be distinguished by the discriminator  $D$ .

$$\min_G \max_D (\mathcal{L}(D, G)) = \min_G \max_D \left( \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_{\text{in}}(z)} [\log(1 - D(G(z)))] \right)$$

During the training, we alternate between  $r_D$  epochs of optimising  $D$  and  $r_G$  of optimising  $G$ . Here, it is essential to choose  $r_D$  in a way that  $G$  changes slowly enough [147].  $r_T$  describes how often alterations are repeated. The authors of [147] choose an algorithm based on *minibatch stochastic gradient*, as shown in Algorithm 2. Further, they recommend using a momentum method for the update. Both of these techniques were explained in Section 2.4.

---

**Algorithm 2** Minibatch stochastic gradient descent for GAN training.

---

**for**  $r_T$  epochs **do**

**for**  $r_D$  epochs **do**

    Sample minibatch  $\{z^i\}_{i=1}^S$  from noise prior  $p_{\text{in}}(z)$ .

    Sample minibatch  $\{x^i\}_{i=1}^S$  from data generating distribution  $p_{\text{data}}(x)$ .

    Update the discriminator by ascending its stochastic gradient

$$\nabla_{\theta_D} \frac{1}{S} \sum_{i=1}^S [\log D(x^i) + \log(1 - D(G(z^i)))]$$

**end for**

**for**  $r_G$  epochs **do**

    Sample minibatch of  $S$  noise samples  $\{z^i, \dots, z^S\}$  from noise prior  $p_{\text{in}}(z)$ .

    Update the generator by descending its stochastic gradient

$$\nabla_{\theta_G} \frac{1}{S} \sum_{i=1}^S \log(1 - D(G(z^i)))$$

**end for**

**end for**

---

At this point, we want to mention one of the main difficulties occurring when training a GAN, called *mode collapse* [349]. This describes the situation where the generative model only produces a selection of the aimed training data. When we think in the example of producing “handwritten” digits, the generator would, for example, produce some of the digits never or only rarely.

Mode collapse often occurs when the discriminator is not trained well enough. On the other hand, when the discriminator is trained optimal, the gradient used in Algorithm 2 vanishes. Different solutions were proposed since this is a common problem in generative ML. One of them is the in [349] discussed *Wasserstein* GAN. Here, instead of the discriminative, a *critic* network is applied. Using the Wasserstein distance [350] the two different kinds of outputs of the critic are

compared and aimed to be maximised by the critic: the output of the critic getting training data as input and the output of it getting the generated data. In the same way as the original GAN, this results in a minimax problem, i.e.

$$\min_G \max_D (\mathcal{L}(D, C)) = \min_G \max_D \left( \frac{1}{S} \sum_{i=1}^S C(x_i) - C(G(z_i)) \right).$$

### Theoretical results

Before discussing GAN applications, we shortly want to mention some theoretical results of the original GAN, proposed in [147], with the aim of getting a better intuition of the training algorithm.

Whereas the generator's and discriminator's data distribution changes during the training, the training samples generating distribution  $p_{\text{data}}$  does not change. In the first phase of Algorithm 2,  $G$  is fixed and only the discriminator is trained. We can formulate this phase as

$$\begin{aligned} \mathcal{L}(G, D) &= \int_x p_{\text{data}}(x) \log(D(x)) dx + \int_z p_{\text{in}}(z) \log(1 - D(G(z))) dz \\ &= \int_x p_{\text{data}}(x) \log(D(x)) + p_G(x) \log(1 - D(x)) dx. \end{aligned}$$

For  $(a, b) \in \mathbb{R}^2 \setminus \{0, 0\}$ , the function  $y \rightarrow a \log(y) + b \log(1 - y)$  achieves its maximum in  $[0, 1]$  at  $\frac{a}{a+b}$ . Hence for a fixed generator  $G$ , the optimal discriminator, denoted by  $D^*$ , is

$$D^*(x) = \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_G(x)}.$$

After training the discriminator, the generator  $G$  gets trained. The authors of [147] observed that the discriminator is more likely to classify the data at first. Experimental results show that after several training iterations, including updating  $D$  and  $G$ , we can reach the state of  $p_G = p_{\text{data}}$ . The discriminator is unable to differentiate between the two distribution, i.e.  $D(x) = \frac{1}{2}$ .

Indeed the authors of [147] showed that the global minimum of the training criterion

$$\begin{aligned} C(G) &= \max_D \mathcal{L}(G, D) \\ &= \mathbb{E}_{x \sim p_{\text{data}}} [\log D_G(x)] + \mathbb{E}_{z \sim p_{\text{in}}} [\log(1 - D_G(G(z)))] \end{aligned}$$

is achieved if and only if  $p_G = p_{\text{data}}$ .

Moreover it was demonstrated that  $p_G$  actually converges to  $p_{\text{data}}$ , given  $G$  and  $D$  have enough capacity at each step of Algorithm 2 so that the discriminator is allowed to reach its optimum given  $G$ , and  $p_G$  is updated so as to improve the criterion

$$\mathbb{E}_{x \sim p_{\text{data}}} [\log D_G^*(x)] + \mathbb{E}_{x \sim p_G} [\log(1 - D_G^*(x))].$$

### Applications

The authors of [147] trained their proposed GANs at a range of datasets, including the already in Section 2.3 discussed MNIST dataset [148]. Therefore we want to take the opportunity to continue the discussion of the demonstrative MNIST dataset and present in Figure 7.2 the input and output of a, in the context of a GAN, trained generator. Although the authors do not claim that the resulting samples are better than samples generated by previously existing methods, the figure shows that after 90000 training epochs, it is possible to recognise digits in the outputs. The output is quite diverse, so no mode collapse accrued.

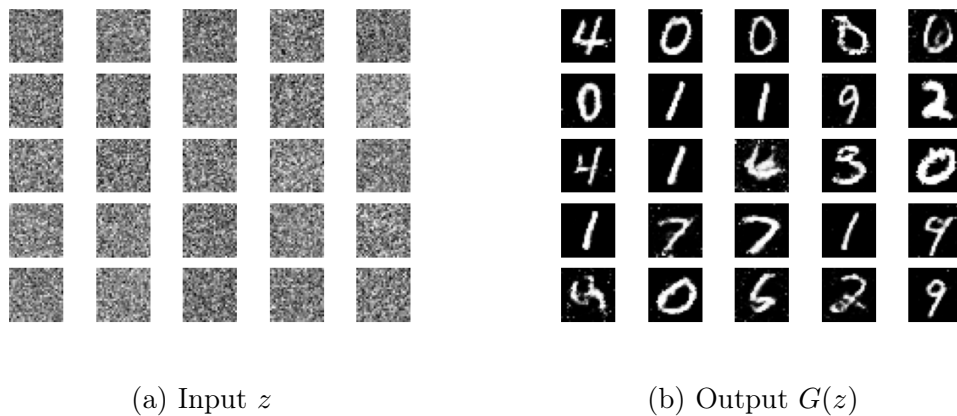


Figure 7.2.: **Generating “handwritten” digits.** Using the code proposed in [147] and the MNIST data set [148] we trained a generator  $G$  in a GAN context. After 90000 epochs the  $G$  generated the output depicted in (b) from the input in (a).

Although generating pictures of digits is indeed a good demonstration example, we will discuss more practicable applications of GANs in the following lines. Based on the original proposal discussed above, generative NNs are applied in many areas today [336]. These network architectures allow training in semi-supervised and unsupervised learning manner, for example, when producing data pairs is expensive or impossible.

So far, we have focused on the output of the generator. However, also the in GAN training acquired discriminator can be of use, for example, for classification or regression tasks [336–338]. Moreover, generators can be used to simply generate more labelled training samples for training further NNs [351, 352].

Also, *image synthesis*, i.e. the generation of images, plays a significant role in the usage of GANs. Many different techniques arose, for example, algorithms where both, the generator and the discriminator, are fed with additional labels. Using the latter ansatz, for example, “reverse captioning”, i.e. the generation of pictures due to their description is possible [339].

In contrast, the MNIST example displayed in Figure 7.2 belongs in the category of GANs used for *image-to-image translation*. Here an input image gets processed to



an output image. The model *pix2pix* [353] can, for example, learn to generate maps from aerial photos or colourise greyscale images. Furthermore, *super-resolution* is a scope of application. Using GANs, low-resolution images can be upscaled in a realistic and natural-looking way [340].

For a more detailed discussion of different kinds of GANs and their applications, we point to [162, 336].

## 7.2 Dissipative quantum generative adversarial network

In analogy to the in Figure 7.1 depicted classical GAN, the in the following presented DQGAN, proposed in [102], is constructed of two DQNNs, the generative model, and the discriminative model, described through the completely positive maps  $\mathcal{E}_G$  and  $\mathcal{E}_D$ , respectively. The number of qubits in the last layer of the generator is equal to the number of qubits in the first layer of the discriminator, so that the generator's output can be used as input for the discriminator. We are given a set of training states  $\{|\phi_x^T\rangle\}_{x=1}^N$ . Moreover, several sets of random states  $\{|\psi_x^{\text{in}}\rangle\}$  are needed during the training algorithm. The aim is that the generative DQNN produces states similar to the training states  $|\phi_x^T\rangle$ . Note that we assume the states of these sets to be pure.




Figure 7.3.: **DQGAN.** The here depicted DQGAN consists of four qubits. Qubits 2 and 3 are shared by the generative and the discriminative DQNN. The state of these qubits is either the generator applied on the input state, i.e.  $\mathcal{E}_G(|\psi^{\text{in}}\rangle \langle\psi^{\text{in}}|)$  or a given training state  $|\phi^T\rangle$ .

Figure 7.3 presents a minimalistic DQGAN: both, the generator and the discriminator consist of two layers and involve three qubits, respectively. Since the qubits of the second layer of the generator, qubits 2 and 3, are shared with the discriminator, the whole DQNN includes four qubits, if the discriminator gets a generated state  $\mathcal{E}_G(|\psi^{\text{in}}\rangle \langle\psi^{\text{in}}|)$  as an input. In the other case, these qubits are directly initialised by a training data state and the DQNN consists only of qubits 2, 3 and 4. In equations we can describe this as

## 7. Quantum generative adversarial networks

$$\rho^{\text{out}} = \begin{cases} \mathcal{E}_D(\mathcal{E}_G(|\psi^{\text{in}}\rangle \langle\psi^{\text{in}}|)) & \text{for generated data} \\ \mathcal{E}_D(|\phi^T\rangle \langle\phi^T|) & \text{for training data.} \end{cases}$$

Note that in DQNN manner, the generator depicted in Figure 7.3 consists of two two-qubit unitaries  $U_{G1}$  and  $U_{G2}$ , acting on qubits 1 and 2, and on qubits 1 and 3. On the contrary, the discriminator is described by a single three-qubit unitary  $U_D$ . To clarify the network structure of the DQGAN, we explain in the following how the states migrate through the network, assuming the depicted example.

The generator gets a random state  $|\psi_x^{\text{in}}\rangle$  as an input. Next, we initialise the output of the generator as  $|00\rangle$ , tensor this state to the input state, apply the network unitaries and trace out the input layer. This can be described as

$$\rho_{G,x} = \mathcal{E}_G(|\psi_x^{\text{in}}\rangle \langle\psi_x^{\text{in}}|) = \text{tr}_{\{1\}} \left( U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle\psi_x^{\text{in}}| \otimes |00\rangle \langle 00|) U_{G1}^\dagger U_{G2}^\dagger \right).$$

The discriminator's input is either the output of the generator  $\rho_{G,x}$  leading to

$$\rho_{G+D,x} = \text{tr}_{\{2,3\}} (U_D(\rho_{G,x} \otimes |0\rangle \langle 0|) U_D^\dagger)$$

or in the other case the training state  $|\phi_x^T\rangle$  leading to

$$\rho_{D,x} = \text{tr}_{\{2,3\}} (U_D(|\phi_x^T\rangle \langle\phi_x^T| \otimes |0\rangle \langle 0|) U_D^\dagger).$$

### 7.3 Loss functions

In the same way as the in Section 4.2 described loss function for the supervised DQNN, the loss functions to train the DQGAN are based on the fidelity. The discriminative model  $\mathcal{E}_D$  aims to output the state  $|1\rangle$  when fed with a state  $|\psi^{\text{in}}\rangle$  and identify “false” data, i.e. states generated by  $\mathcal{E}_G$ , with outputting the state  $|0\rangle$ .

During the training only  $S < N$  of the states  $\{|\phi_x^T\rangle\}_{x=1}^N$  are used. However, we aim that at the end of the training that every of the generator's output states is close to a state of the full data set  $\{|\phi_x^T\rangle\}_{x=1}^N$ . In other words we desire that the generator does not reproduces only the  $S$  given states but extends this data set.

#### Training loss

In analogy to the classical case we can describe the training process through

$$\min_G \max_D \left( \frac{1}{S} \sum_{x=1}^S \langle 0| \mathcal{E}_D(\mathcal{E}_G(|\psi_x^{\text{in}}\rangle \langle\psi_x^{\text{in}}|)) |0\rangle + \frac{1}{S} \sum_{x=1}^S \langle 1| \mathcal{E}_D(|\phi_x^T\rangle \langle\phi_x^T|) |1\rangle \right), \quad (7.1)$$

and the updates of the discriminator and the generator take place consecutively. For updating the generator we maximise the loss function

$$\mathcal{L}_D(\mathcal{E}_D, \mathcal{E}_G) = \frac{1}{S} \sum_{x=1}^S \langle 0 | \mathcal{E}_D(\mathcal{E}_G(|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}|)) | 0 \rangle + \frac{1}{S} \sum_{x=1}^S \langle 1 | \mathcal{E}_D(|\phi_x^T\rangle \langle \phi_x^T|) | 1 \rangle.$$

The generator is trained through maximising

$$\mathcal{L}_G(\mathcal{E}_D, \mathcal{E}_G) = \frac{1}{S} \sum_{x=1}^S \langle 1 | \mathcal{E}_D(\mathcal{E}_G(|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}|)) | 1 \rangle.$$

As shown in Algorithm 3, in every training round two new sets of random states  $\{|\psi_x^{\text{in}}\rangle\}_{x=1}^S$  are used to evaluate  $\mathcal{L}_D$  and  $\mathcal{L}_G$ , respectively. Note further that  $\mathcal{L}_G$  differs from the corresponding term in Equation (7.1): the fidelity is calculated with respect to  $|1\rangle$  instead of  $|0\rangle$  and hence the generator is trained by maximising  $\mathcal{L}_G$  rather than minimising. Using this arrangement it will be more convenient to compare the loss functions in Section 7.5 and Section 7.6.

### Validation loss

In the final stages of the training, we aim that every output of the generator is close to one of the given states  $\{|\phi_x^T\rangle\}_{x=1}^N$ . Therefore we produce  $V$  additional random states  $\{|\psi_x^{\text{in}}\rangle\}_{x=1}^V$ . For every of these validation states, we search for the best suiting state of the data set through  $\max_{x=1}^N (\langle \phi_x^T | \mathcal{E}_G(|\psi_i^{\text{in}}\rangle \langle \psi_i^{\text{in}}|) | \phi_x^T \rangle)$  and build the mean value over all  $V$  generated states, i.e.

$$\mathcal{L}_V(\mathcal{E}_D, \mathcal{E}_G) = \frac{1}{V} \sum_{i=1}^V \max_{x=1}^N (\langle \phi_x^T | \mathcal{E}_G(|\psi_i^{\text{in}}\rangle \langle \psi_i^{\text{in}}|) | \phi_x^T \rangle).$$

This *validation loss* function is used to test the training success.

## 7.4 Training algorithm

Since the presented algorithm consists of two parts, we define the training epochs with three parameters: the overall repetition of alternating between training the discriminator and the generator  $r_T$ , a parameter  $r_D$  describing the repetition of training discriminator and a parameter  $r_G$  describing the repetition of training generator. These parameters are used in the *for* loops of the pseudo-code in Algorithm 3, which gives an overview of the algorithm. The last three steps describe the validation process after the training.

**Algorithm 3** Training of a DQGAN.

---

```

initialise network unitaries
for  $r_T$  epochs do
  make a list of  $S$  randomly chosen states of the training data list  $\{|\phi_x^T\rangle\}_{x=1}^N$ 
  for  $r_D$  epochs do
    make a list of  $S$  random states  $|\psi_x^{\text{in}}\rangle$ 
    update the discriminator unitaries with maximising  $\mathcal{L}_D$ 
  end for
  for  $r_G$  epochs do
    make a list of  $S$  random states  $|\psi_x^{\text{in}}\rangle$ 
    update the generator unitaries with maximising  $\mathcal{L}_G$ 
  end for
end for
make a list of  $V$  random states  $|\psi_x^{\text{in}}\rangle$ 
propagate each  $|\psi_x^{\text{in}}\rangle$  through the generator to produce  $V$  new states
compute  $\mathcal{L}_V$ 

```

---

**Derivation of the update matrices**

Analogous to the DQNN update rule presented in Section 4.3 the unitaries will be updated through

$$U_j^l(t + \epsilon) = e^{i\epsilon K_j^l(t)} U_j^l(t).$$

Before we derive the form of the update matrices  $K_j^l$  in general we take a smaller step first and discuss the update of the minimal example which is depicted in Figure 7.3 and consists of three unitaries. These unitaries can be updated via

$$\begin{aligned} U_D(t + \epsilon) &= e^{i\epsilon K_D(t)} U_D(t) \\ U_{G1}(t + \epsilon) &= e^{i\epsilon K_{G1}(t)} U_{G1}(t) \\ U_{G2}(t + \epsilon) &= e^{i\epsilon K_{G2}(t)} U_{G2}(t). \end{aligned}$$

We want to use this minimal example and explain how the update changes the output states. Note that in the following the unitaries act on the current layers, e.g.  $U_{G1}$  is actually  $U_{G1} \otimes \mathbb{1}$  and  $U_{G2}$  represents  $\mathbb{1} \otimes U_{G2}$ .

In the first part of the algorithm the generator is fixed and the discriminator is updated. When the discriminator is fed with training data we get the output state

$$\begin{aligned} \rho_{\text{out}}^D(t + \epsilon) &= \text{tr}_{\{2,3\}} \left( e^{i\epsilon K_D} U_D (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\rangle \langle 0|) U_D^\dagger e^{-i\epsilon K_D} \right) \\ &= \text{tr}_{\{2,3\}} \left( U_D (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\rangle \langle 0|) U_D^\dagger + i\epsilon \left[ K_D, U_D (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\rangle \langle 0|) U_D^\dagger \right] \right. \\ &\quad \left. + \mathcal{O}(\epsilon^2) \right) \\ &= \rho_{\text{out}}^D(t) + i\epsilon \text{tr}_{\{2,3\}} \left( \left[ K_D, U_D (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\rangle \langle 0|) U_D^\dagger \right] \right) + \mathcal{O}(\epsilon^2). \end{aligned}$$

In the case the discriminator gets input of the generator we get the output state

$$\begin{aligned}
 \rho_{\text{out}}^{G+D}(t + \epsilon) &= \text{tr}_{\{1,2,3\}} \left( e^{i\epsilon K_D} U_D U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger U_D^\dagger e^{-i\epsilon K_D} \right) \\
 &= \text{tr}_{\{1,2,3\}} \left( U_D U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger U_D^\dagger \right. \\
 &\quad \left. + i\epsilon \left[ K_D, U_D U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger U_D^\dagger \right] + \mathcal{O}(\epsilon^2) \right) \\
 &= \rho_{\text{out}}^{G+D}(t) \\
 &\quad + i\epsilon \text{tr}_{\{1,2,3\}} \left( \left[ K_D, U_D U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger U_D^\dagger \right] \right) \\
 &\quad + \mathcal{O}(\epsilon^2).
 \end{aligned}$$

The update of the generator, assuming a fixed discriminator, is

$$\begin{aligned}
 \rho_{\text{out}2}^{G+D}(t + \epsilon) &= \text{tr}_{\{1,2,3\}} \left( U_D e^{i\epsilon K_{G2}} U_{G2} e^{i\epsilon K_{G1}} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) \right. \\
 &\quad \left. U_{G1}^\dagger e^{-i\epsilon K_{G1}} U_{G2}^\dagger e^{-i\epsilon K_{G2}} U_D^\dagger \right) \\
 &= \text{tr}_{\{1,2,3\}} \left( U_D \left( U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger \right. \right. \\
 &\quad \left. \left. + i\epsilon U_{G2} \left[ K_{G1}, U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger \right] U_{G2}^\dagger \right. \right. \\
 &\quad \left. \left. + i\epsilon \left[ K_{G2}, U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger \right] \right) U_D^\dagger + \mathcal{O}(\epsilon^2) \right) \\
 &= \rho_{\text{out}2}^{G+D}(t) \\
 &\quad + i\epsilon \text{tr}_{\{1,2,3\}} \left( U_D \left( U_{G2} \left[ K_{G1}, U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger \right] U_{G2}^\dagger \right. \right. \\
 &\quad \left. \left. + \left[ K_{G2}, U_{G2} U_{G1} (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |000\rangle \langle 000|) U_{G1}^\dagger U_{G2}^\dagger \right] \right) U_D^\dagger \right) + \mathcal{O}(\epsilon^2).
 \end{aligned}$$

To derive the update matrices in general, i.e. beyond this minimal example, we assume in the following a generator consisting of unitaries  $U_1^1 \dots U_{m_g}^g$  and a discriminator built of unitaries  $U_1^{g+1} \dots U_{m_{L+1}}^{L+1}$ . The update matrices  $K_j^l$  update the generator if  $l \leq g$ , where  $g$  is the generator's number of perceptron layers. Otherwise, the matrices describe discriminator updates. Note that in the sections describing the numerical results, Section 7.5 and Section 7.6, we will always use  $g = 1$ , i.e. a generator with just one perceptron layer connecting two layers of qubits.

**Proposition 7.1.** *The update matrix for a DQGAN trained with pure states  $|\phi_x^T\rangle$  has to be of the form*

$$K_j^l(t) = \frac{\eta 2^{m_l - 1} i}{S} \sum_x \text{tr}_{\text{rest}} (M_j^l(x, t)),$$

## 7. Quantum generative adversarial networks

where

$$M_j^l = \left[ U_j^l \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_j^{l\dagger}, \right. \\ \left. U_{j+1}^{l\dagger} \dots U_{m_{L+1}}^{L+1\dagger} (\mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1|) U_{m_{L+1}}^{L+1} \dots U_{l+1}^l \right]$$

for  $l \leq g$  and

$$M_j^l = \left[ U_j^l \dots U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{g+1\dagger} \dots U_j^{l\dagger} \right. \\ \left. - U_j^l \dots U_1^{g+1} U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_g}^{g\dagger} U_1^{g+1\dagger} \dots U_j^{l\dagger}, \right. \\ \left. U_{j+1}^{l\dagger} \dots U_{m_{L+1}}^{L+1\dagger} (\mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1|) U_{m_{L+1}}^{L+1} \dots U_{l+1}^l \right]$$

else. Here  $U_j^l$  is assigned to the  $j$ th perceptron acting on layers  $l-1$  and  $l$ ,  $g$  is the number of perceptron layers of the generator, and  $\eta$  is the learning rate.

*Proof.* To study the training of the discriminator, we need to compute the output state after an update with  $K_D$ . Note that in the following the unitaries act on the current layers, e.g.  $U_1^l$  denotes actually  $U_1^l \otimes \mathbb{1}_{2,3,\dots,m_l}^l$ . First we fix the generator. To derive the update for the discriminator, we compute the state when it is fed with the training data, i.e.

$$\rho_{\text{out}}^D(t + \epsilon) = \text{tr}_{\text{in(D)+hid}} \left( e^{i\epsilon K_{m_{L+1}}^{L+1}} U_{m_{L+1}}^{L+1} \dots e^{i\epsilon K_1^{g+1}} U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) \right. \\ \left. U_1^{g+1\dagger} e^{-i\epsilon K_1^{g+1}} \dots U_{m_{L+1}}^{L+1\dagger} e^{-i\epsilon K_{m_{L+1}}^{L+1}} \right) \\ = \rho_{\text{out}}^D(t) + i\epsilon \text{tr}_{\text{in(D)+hid}} \left( [K_{m_{L+1}}^{L+1}, U_{m_{L+1}}^{L+1} \dots U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) \right. \\ \left. U_1^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger}] + \dots + U_{m_{L+1}}^{L+1} \dots U_2^{g+1} [K_1^{g+1}, U_1^{g+1} |\phi_x^T\rangle \langle \phi_x^T| \right. \\ \left. \otimes |0\dots 0\rangle \langle 0\dots 0| U_1^{g+1\dagger}] U_2^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) + \mathcal{O}(\epsilon^2).$$

Analogously we formulate the state when the discriminator gets the generator's output as input as

$$\rho_{\text{out}}^{G+D}(t + \epsilon) = \text{tr}_{\text{in(G)+hid}} \left( e^{i\epsilon K_{m_{L+1}}^{L+1}} U_{m_{L+1}}^{L+1} \dots e^{i\epsilon K_1^{g+1}} U_1^{g+1} U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \right. \\ \left. \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_g}^{g\dagger} U_1^{g+1\dagger} e^{-i\epsilon K_1^{g+1}} \dots U_{m_{L+1}}^{L+1\dagger} e^{-i\epsilon K_{m_{L+1}}^{L+1}} \right) \\ = \rho_{\text{out}}^D(t) + i\epsilon \text{tr}_{\text{in(G)+hid}} \left( [K_{m_{L+1}}^{L+1}, U_{m_{L+1}}^{L+1} \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \right. \\ \left. \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_{L+1}}^{L+1\dagger}] + \dots \right. \\ \left. + U_{m_{L+1}}^{L+1} \dots U_2^{g+1} [K_1^{g+1}, U_1^{g+1} U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) \right. \\ \left. U_1^{1\dagger} \dots U_{m_g}^{g\dagger} U_1^{g+1\dagger}] U_2^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) + \mathcal{O}(\epsilon^2).$$

Further, the derivative of the discriminator loss function is of the form

$$\begin{aligned}
\frac{d\mathcal{L}_D}{dt} &= \lim_{\epsilon \rightarrow 0} \frac{\mathcal{L}_D(t) + i\epsilon \frac{1}{S} \sum_{x=1}^S \langle 1 | \text{tr}_{\text{in+hid}}(\dots) | 1 \rangle - \mathcal{L}_D(t)}{\epsilon} \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr}_{\text{in+hid}} \left( \mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1| \left( \left( [K_{m_{L+1}}^{L+1}, U_{m_{L+1}}^{L+1}} \dots U_1^{g+1} |\phi_x^T\rangle \langle \phi_x^T| \right. \right. \right. \\
&\quad \otimes |0\dots 0\rangle \langle 0\dots 0| U_1^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger}] + \dots \\
&\quad + U_{m_{L+1}}^{L+1} \dots U_2^{g+1} [K_1^{g+1}, U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{g+1\dagger}] \\
&\quad \left. \left. \left. U_2^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) \right. \right. \\
&\quad - \left( [K_{m_{L+1}}^{L+1}, U_{m_{L+1}}^{L+1}} \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_{L+1}}^{L+1\dagger}] + \dots \right. \\
&\quad \left. \left. \left. + U_{m_{L+1}}^{L+1} \dots U_2^{g+1} [K_1^{g+1}, U_1^{g+1} U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) \right. \right. \right. \\
&\quad \left. \left. \left. U_1^{1\dagger} \dots U_{m_g}^{g\dagger} U_1^{g+1\dagger}] U_2^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) \right) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr}_{\text{in+hid}} \left( \left[ U_{m_{L+1}}^{L+1} \dots U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right. \right. \\
&\quad \left. \left. - U_{m_{L+1}}^{L+1} \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_{L+1}}^{L+1\dagger}, \right. \right. \\
&\quad \left. \left. \mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1| \right] K_{m_{L+1}}^{L+1} + \dots + \left[ U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{g+1\dagger} \right. \right. \\
&\quad \left. \left. - U_1^{g+1} U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) \right. \right. \\
&\quad \left. \left. U_1^{1\dagger} \dots U_{m_g}^{g\dagger} U_1^{g+1\dagger}, U_{m_{L+1}}^{L+1\dagger} \dots U_2^{g+1\dagger} \mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1| U_2^{g+1} \dots U_{m_{L+1}}^{L+1} \right] K_1^{g+1} \right) \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr}_{\text{in+hid}} \left( M_{m_{L+1}}^{L+1} K_{m_{L+1}}^{L+1} + \dots + M_1^{g+1} K_1^{g+1} \right).
\end{aligned}$$

Note that at this point  $|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|$  denotes  $\mathbb{1}_{\text{in}(G)+\text{hid}(G)} \otimes |\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|$ , to match the dimension of the other summand. To this point the generator was fixed. As a next step we fix the discriminator. Using the state

$$\begin{aligned}
\rho_{\text{out}2}^{G+D}(s + \epsilon) &= \text{tr}_{\text{in}(G)+\text{hid}} \left( U_{m_{L+1}}^{L+1} \dots U_1^{g+1} e^{i\epsilon K_{m_g}^g} U_{m_g}^g \dots e^{i\epsilon K_1^1} U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \right. \\
&\quad \left. \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} e^{-i\epsilon K_1^1} \dots U_{m_g}^{g\dagger} e^{-i\epsilon K_{m_g}^g} U_1^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) \\
&= \rho_{\text{out}}^D(t) + i\epsilon \text{tr}_{\text{in}(G)+\text{hid}} \left( U_{m_{L+1}}^{L+1} \dots U_1^{g+1} [K_{m_g}^g, U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \right. \\
&\quad \left. \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_g}^{g\dagger}] U_1^{g+1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} + \dots \right. \\
&\quad \left. + U_{m_{L+1}}^{L+1} \dots U_2^1 [K_1^1, U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger}] U_2^{1\dagger} \dots U_{m_{L+1}}^{L+1\dagger} \right) \\
&\quad + \mathcal{O}(\epsilon^2)
\end{aligned}$$

## 7. Quantum generative adversarial networks

the derivative of the loss function for training the generator can be written as

$$\begin{aligned}
\frac{d\mathcal{L}_G}{dt} &= \lim_{\epsilon \rightarrow 0} \frac{\mathcal{L}_G(t) + i\epsilon \frac{1}{S} \sum_x \langle 1 | \text{tr}_{\text{in+hid}}(\dots) | 1 \rangle - \mathcal{L}_G(t)}{\epsilon} \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr} \left( \mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1| \left( \left( U_{m_{l+1}}^{l+1} \dots U_1^{g+1} [K_{m_g}^g, U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \right. \right. \right. \\
&\quad \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_g}^{g\dagger}] U_1^{g+1\dagger} \dots U_{m_{l+1}}^{l+1\dagger} + \dots \\
&\quad \left. \left. \left. + U_{m_{l+1}}^{l+1} \dots U_2^1 [K_1^1, U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger}] U_2^{1\dagger} \dots U_{m_{l+1}}^{l+1\dagger} \right) \right) \right) \\
&= \frac{i}{S} \sum_{x=1}^S \text{tr} \left( \left[ U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_g}^{g\dagger}, \right. \right. \\
&\quad \left. U_1^{g+1\dagger} \dots U_{m_{l+1}}^{l+1\dagger} (\mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1|) U_{m_{l+1}}^{l+1} \dots U_1^{g+1} \right] K_{m_g}^g + \dots \\
&\quad + \left[ U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger}, \right. \\
&\quad \left. U_2^{1\dagger} \dots U_{m_{l+1}}^{l+1\dagger} (\mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1|) U_{m_{l+1}}^{l+1} \dots U_2^1 \right] K_1^1 \Big) \\
&\equiv \frac{i}{S} \sum_{x=1}^S \text{tr} \left( M_{m_g}^g K_{m_g}^g + \dots + M_1^1 K_1^1 \right).
\end{aligned}$$

In both updates, we parametrise the parameter matrices analogously to the proof of Proposition 4.1 as

$$K_j^l(t) = \sum_{\alpha_1, \alpha_2, \dots, \alpha_{m_{l-1}}, \beta} K_{j, \alpha_1, \dots, \alpha_{m_{l-1}}, \beta}^l(t) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^{\alpha_{m_{l-1}}} \otimes \sigma^\beta),$$

where the  $\alpha_i$  denote the qubits in the previous layer and  $\beta$  denotes the current qubit in layer  $l$ . To reach the maximum of the loss function as a function of the parameters *fastest*, we maximise  $\frac{d\mathcal{L}}{dt}$ . Since this is a linear function, the extrema are at  $\pm\infty$ . To ensure that we get a finite solution, we introduce a Lagrange multiplier  $\lambda \in \mathbb{R}$ . Hence, to find  $K_j^l$  we have to solve a following maximisation problem for the discriminator update and for the generator update. Since both are solved analogously we only formulate the maximisation for the discriminator update here, namely

$$\begin{aligned}
&\max_{K_{j, \alpha_1, \dots, \beta}^l} \left( \frac{dC(t)}{dt} - \lambda \sum_{\alpha_i, \beta} K_{j, \alpha_1, \dots, \beta}^l(t)^2 \right) \\
&= \max_{K_{j, \alpha_1, \dots, \beta}^l} \left( \frac{i}{S} \sum_{x=1}^S \text{tr} \left( M_{m_{L+1}}^{L+1} K_{m_{L+1}}^{L+1} + \dots + M_1^{g+1} K_1^{g+1} \right) - \lambda \sum_{\alpha_1, \dots, \beta} K_{j, \alpha_1, \dots, \beta}^l(t)^2 \right) \\
&= \max_{K_{j, \alpha_1, \dots, \beta}^l} \left( \frac{i}{S} \sum_{x=1}^S \text{tr}_{\alpha_1, \dots, \beta} \left( \text{tr}_{\text{rest}} \left( M_{m_{L+1}}^{L+1} K_{m_{L+1}}^{L+1} + \dots + M_1^{g+1} K_1^{g+1} \right) \right) \right)
\end{aligned}$$



$$- \lambda \sum_{\alpha_1, \dots, \beta} K_{j, \alpha_1, \dots, \beta}^l(t)^2).$$

Taking the derivative with respect to  $K_{j, \alpha_1, \dots, \beta}^l$  leads to

$$\frac{i}{S} \sum_{x=1}^S \text{tr}_{\alpha_1, \dots, \beta} (\text{tr}_{\text{rest}} (M_j^l(t)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta)) - 2\lambda K_{j, \alpha_1, \dots, \beta}^l(t) = 0,$$

hence,

$$K_{j, \alpha_1, \dots, \beta}^l(t) = \frac{i}{2S\lambda} \sum_{x=1}^S \text{tr}_{\alpha_1, \dots, \beta} (\text{tr}_{\text{rest}} (M_j^l(t)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta))$$

This produces the matrix

$$\begin{aligned} K_j^l(t) &= \sum_{\alpha_1, \dots, \beta} K_{j, \alpha_1, \dots, \beta}^l(t) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \\ &= \frac{i}{2S\lambda} \sum_{\alpha_1, \dots, \beta} \sum_{x=1}^S \text{tr}_{\alpha_1, \dots, \beta} (\text{tr}_{\text{rest}} (M_j^l(t)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta)) (\sigma^{\alpha_1} \otimes \dots \otimes \sigma^\beta) \\ &= \frac{\eta 2^{m_l-1} i}{2S} \sum_{x=1}^S \text{tr}_{\text{rest}} (M_j^l(t)), \end{aligned}$$

where  $\eta = 1/\lambda$  is the learning rate and  $\text{tr}_{\text{rest}}$  traces out all qubits that the perceptron unitary  $U_j^l$  does not act on.

Notice again that  $K_j^l$  updates the generator if  $j \leq g$  for the generator's number of perceptron layers  $g$ . The definition of  $M_j^l$  is

$$\begin{aligned} M_j^l &= \left[ U_j^l \dots U_1^l (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_j^{l\dagger}, \right. \\ &\quad \left. U_{j+1}^{l\dagger} \dots U_{m_{L+1}}^{L+1\dagger} (\mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1|) U_{m_{L+1}}^{L+1} \dots U_{l+1}^l \right] \end{aligned}$$

for  $l \leq g$  and

$$\begin{aligned} M_j^l &= \left[ U_j^l \dots U_1^{g+1} (|\phi_x^T\rangle \langle \phi_x^T| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{g+1\dagger} \dots U_j^{l\dagger} \right. \\ &\quad \left. - U_j^l \dots U_1^{g+1} U_{m_g}^g \dots U_1^1 (|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}| \otimes |0\dots 0\rangle \langle 0\dots 0|) U_1^{1\dagger} \dots U_{m_g}^{g\dagger} U_1^{g+1\dagger} \dots U_j^{l\dagger}, \right. \\ &\quad \left. U_{j+1}^{l\dagger} \dots U_{m_{L+1}}^{L+1\dagger} (\mathbb{1}_{\text{in+hid}} \otimes |1\rangle \langle 1|) U_{m_{L+1}}^{L+1} \dots U_{l+1}^l \right] \end{aligned}$$

else. □

## 7.5 Classical simulation

Just like in Chapter 4 and Chapter 6, where we discussed the initially proposed DQNNs [79] and the extension to graph-structured quantum data [99], we use classical simulation to study the above-described training methods for DQGANs.

In contrary to the corresponding sections in the chapters mentioned above, we structure this section into two parts. The first part describes how the data set used for training the DQGAN. Next, we use the classical simulation to demonstrate how the generator’s and discriminator’s training losses and the validation loss evolve during the training process. In the last part, we explain how to study the generated states’ diversity. The used code can be found at [276].

### Training set

As an exemplary data set we use a set of pure one-qubit states which build a line on the Bloch sphere, namely

$$\text{data}_{\text{line}} = \left\{ \frac{(N-x)|0\rangle + (x-1)|1\rangle}{\|((N-x)|0\rangle + (x-1)|1\rangle)\|} \right\}_{x=1}^N, \quad (7.2)$$


see also Figure 7.5a. To use  $\text{data}_{\text{line}}$  for the training it is shuffled. The first  $S$  of the resulting set, denoted as  $\{|\phi_x^T\rangle\}_{x=1}^S$ , will be used for the training process and can be accessed by the generative model. The full data set  $\{|\phi_x^T\rangle\}_{x=1}^N$  with  $N > S$  is used for computing the validation loss. In the following we choose  $N = 50$  and  $S = 10$ .

### Evaluation of the loss functions

In Figure 7.4a, both training losses and the validation loss during training of a 1-1-1 DQGAN are plotted. For all in the following presented examples we choose  $r_D = r_G = 1$ ,  $\epsilon = 0.01$  and  $\eta = 1$  and a network architecture of 1-1-1. Note that in Figure C.1 in the appendix, the training of a 1-3-1 DQGAN is documented as well.

Compared to the prior loss function plots in this work, the evaluation of the losses is much more eventful. The validation loss reaches values over 0.95 after  $r_T = 475$  training epochs. In the first training epochs, the training loss of the generator shrinks and the discriminator training loss increases. This behaviour is inverted after  $r_T = 100$ . Following the remaining training, we find that this process is repeated. The generator’s and discriminator’s opposing goals, formulated in the training loss functions, explain this behaviour.

Furthermore, we can observe the saturation of the validation loss at nearly 1. This means that every output of the generator is close to one state of  $\text{data}_{\text{line}}$ . Note that here we define the closeness through the fidelity of two quantum states: the larger the fidelity is, the closer the states are. However, from this plot, it is not clear if the generator produces the same state time after time or, in the other extreme, reproduces all states of the set  $\text{data}_{\text{line}}$ . In both cases, the validation



(a) Loss functions.




(b) Diversity of the generator's output after  $r_T = 300$  training epochs.(c) Diversity of the generator's output after  $r_T = 500$  training epochs.(d) Diversity of the generator's output after  $r_T = 800$  training epochs.

Figure 7.4.: **Training a DQGAN.** (a) depicts the evolution of the loss functions during the training of a  $\circ-\circ-\circ$  DQGAN in  $r_T = 1000$  epochs with  $\eta = 1$  and  $\epsilon = 0.01$  using 50 data pairs where 10 are used as training pairs. The dashed lines mark the diversity checks after 300 (b), 500 (c) and 800 (d) training epochs  $r_T$  for the generator's output.

## 7. Quantum generative adversarial networks

loss would be maximal. Hence, we have to study the *diversity* of the generator's output in the following.

### Diversity of the generator's output

Since our aim is to train the generator for producing extended data sets, we want its output to be diverse. We build a set of 100 by the generator produced states to test the diversity. For each of these, we find the element with the index  $x$  in  $\text{data}_{\text{line}}$ , which is the closest concerning the fidelity of two states.

Following this procedure, we count how many times every element in  $\text{data}_{\text{line}}$  was (approximately) reproduced by the generator. In Figure 7.4b to 7.4d these numbers are depicted in the form of histograms. According to the definition of  $\text{data}_{\text{line}}$  in Equation (7.2),  $x = 1$  corresponds to the state  $|0\rangle$  and  $x = 50$  to  $|1\rangle$ , respectively. The number in the histogram related to, for example,  $x = 1$  shows how often the output of the generator was approximately  $|0\rangle$ . The different colours describe whether an element of  $\text{data}_{\text{line}}$  was used as a training state or not. We can see that the DQGAN can generalise, and not only the training states get reproduced. Specifically, we show this diversity checks after 300, 500, and 800 training epochs. Comparing Figure 7.4b to 7.4d with the plots of the loss functions in Figure 7.4a (see the dashed lines), let us resume that the training of the DQGAN has to be stopped, in this case, after about 300 training epochs to have a good balance between a large diversity and a significant validation loss.

Figure 7.5, where the output of the generator is plotted for different training epochs in Bloch spheres, confirms this statement. At every of these training steps we build a set of 100 states produced by the generator and plot these states in a Bloch sphere. For the discussion of other training data sets we point to Figure C.2 in the appendix.




Figure 7.5.: **Output of the generator.** To compare the output of the generator (b-i), during the training of a  $\circ\text{-}\circ\text{-}\circ$  DQGAN, to the data set  $\text{data}_{\text{line}}$  (a) we plot the states in Bloch spheres.

## 7.6 NISQ device implementation

The last section described results of the classical simulation of the DQGAN. In the following we want to discuss  $\text{DQGAN}_{\text{NISQ}}$ , i.e. the implementation of the DQGAN algorithm for a NISQ device. In Section 4.6 we described  $\text{DQNN}_{\text{NISQ}}$  in full detail.

Two main things are different in the implementation of  $\text{DQGAN}_{\text{NISQ}}$  compared to the  $\text{DQGAN}_{\text{NISQ}}$ . Firstly, we have two instead of one training functions. Depending on the phase of the algorithm, see Algorithm 3, we update the parameters  $\omega^G_s$  of the generator or the discriminator's parameters  $\omega^D_s$  using the gradient of the training loss  $\mathcal{L}_G$  or  $\mathcal{L}_D$ , respectively. Whereas one measurement per input is needed for computing  $\mathcal{L}_G$ , two are needed for executing  $\mathcal{L}_D$  for one input.

Secondly, we have to implement the quantum circuit, in order that the discriminator can be used in two ways, i.e.

$$\rho^{\text{out}} = \begin{cases} \mathcal{E}_D(\mathcal{E}_G(|\psi^{\text{in}}\rangle \langle \psi^{\text{in}}|)) & \text{for generated data} \\ \mathcal{E}_D(|\phi^T\rangle \langle \phi^T|) & \text{for training data.} \end{cases}$$

Moreover, we could observe that a small change compared to the  $\text{DQNN}_{\text{NISQ}}$  implementation presented in Section 4.6 is helpful for the training the  $\text{DQGAN}_{\text{NISQ}}$ . We describe the modification, which increases the computational power of the  $\text{DQGAN}_{\text{NISQ}}$  without using additional qubits, in the appendix in Figure C.3. In the following we denote this modification with a  $+$  when describing the  $\text{DQNN}$  architecture, for example in 2-3-2 $^+$ . Further we use different learning rates for the generator and discriminator, denoted by  $\eta_D$  and  $\eta_G$ , respectively.

For training a  $\text{DQGAN}_{\text{NISQ}}$  we use the same loss functions  $\mathcal{L}_G$  and  $\mathcal{L}_D$  as well as the same training data,  $\text{data}_{\text{line}}$ . We choose  $S = 10$  equally spaced training states of the set for the training. Further, we select that for each of the  $r_T$  epochs, the discriminator is trained  $r_D = 4$  times with a learning rate  $\eta_D = 0.5$  and the generator  $r_G = 1$  times with a learning rate  $\eta_G = 0.1$ . The results of the training are depicted in Figure 7.6.

After  $r_T = 100$  epochs, see Figure 7.6a, we see that the generator produces states in a little more than half of the training data range. The generator's diversity is improved to two-thirds of the training data range after  $r_T = 440$  training epochs which is depicted in Figure 7.6b. Although not all states of the aimed data sets get produced by the generator, a majority of the produced states is closer to a validation state than a training state. We see this as a training success, since the generator does not only learn to reproduce the training states but instead learns to extend the given training data. For more numerical results using  $\text{DQGAN}_{\text{NISQ}}$  we point to Chapter C and [354].



(a) Diversity of the generator's output after  $r_T = 100$  training epochs.(b) Diversity of the generator's output after  $r_T = 440$  training epochs.

Figure 7.6.: **Training a DQGAN<sub>NISQ</sub>**. The training set features  $S = 10$  equally spaced quantum states from  $\text{data}_{\text{line}}$ . The remaining states from  $\text{data}_{\text{line}}$  are used as validation states. The DQGAN<sub>NISQ</sub> features a 1-1<sup>+</sup> generator and a 1-1<sup>+</sup> discriminator, and employs the hyper-parameters  $r_D = 4$ ,  $\eta_D = 0.5$ ,  $r_G = 1$  and  $\eta_G = 0.1$ .

The numerical experiments with DQGANs and DQGAN<sub>NISQ</sub> have shown that DQNNs can be trained in a generative adversarial context. Moreover, we could observe that, in contrast to the DQNN and graph-DQNN, the choice of the number of training rounds is crucial: a DQGAN trained in too many training epochs leads to low diversity in the generator's output states. This result builds almost the end of this thesis: in the ensuing chapter, we want to conclude the results of this work and give an outlook.





# 8

---

## Conclusion and outlook

---

In this thesis, we investigated *quantum neural networks* (QNNs). These architectures combine two of the most exciting research areas of the 21st century: machine learning and quantum computation. We summarise our results and name ideas for future research questions in the following.

We entered the topic of QNNs in two steps. Firstly, we introduced the reader to their classical counterparts. There, we began with explaining the functionality of the building blocks of *neural networks* (NNs) and their activation functions. It followed an overview of some popular network architectures and methods. Our focus was on supervised feed-forward NNs. Next, the general approach of optimising NNs was clarified. With a view to the following chapters, we focused on the gradient descent and back-propagation methods.

Secondly, we gave an overview of the field of *quantum information*. We introduced *qubits*, the quantum analogous to the classical binary bits, and explained their characteristics, including the phenomenons of superposition, mixed states and entanglement. To prepare for the definitions of quantum loss functions, we explained how to compare two quantum states. Moreover, we introduced quantum circuits, which are used to exploit quantum mechanics for quantum computing, and gave an intuition on how the characteristics of quantum mechanics can be used to outperform classical algorithms. Since quantum computation became experimentally possible in the last years, we described the state of the art of these devices. Moreover, we introduced the topic of QNNs with the focus on the implementation, challenges and opportunities.

### Dissipative quantum neural networks

We presented *dissipative quantum neural networks* (DQNNs) [79, 99, 101, 102], which are designed for fully quantum learning tasks, and are capable of universal quantum computation. These types of QNNs are built of layers of qubits, which are connected via perceptrons. Such a perceptron is engineered as an arbitrary unitary operation and acts on qubits of two consecutive qubit layers.

We described the propagation of an input state  $\rho^{\text{in}}$  through the network using a completely positive map  $\mathcal{E}$ . This map is defined as a composition of layer-to-layer transition maps, i.e.  $\rho^{\text{out}} = \mathcal{E}(\rho^{\text{in}})$ . Such a transition map not only contains tensoring the current layer's state to the state of the next layer qubits and applying the perceptron unitaries operations, but also tracing out the qubits from the first of the two layers. For this reason, these QNNs are called *dissipative*.

Further, we presented training and validation loss functions, which are based on the fidelity of two states and compare the output state of the network with the desired state. Moreover, we discussed how the perceptron unitaries get updated in every training epoch based on the knowledge gained from the computation of the training loss. We formulated the update of the unitary  $U_j^l$ , assigned to the  $j$ th perceptron acting on layers  $l-1$  and  $l$ , as  $U_j^l \rightarrow e^{ieK_j^l}U_j^l$ . Next, we made clear that for such an update, only two states are needed: the output state of the previous layer,  $\rho^{l-1}$ , obtained by feed-forward propagation through the network, and the state of the following layer  $\sigma^l$ , obtained by back-propagation of the desired output up to the current layer. This allows us to train deep DQNNs since the memory requirements scale only with the width, not the length of the QNN.

After explaining why DQNNs are capable of universal quantum computing, we presented examples of classical simulations. We could observe that using  $S$  training data pairs of the form  $\{|\phi^{\text{in}}\rangle, Y|\phi^{\text{in}}\rangle\}$ , these QNNs can learn the unknown unitary  $Y$  in a feasible number of training epochs, which could be demonstrated by using validation data pairs of the same form. Varying the number  $S$  has shown that for only a small number of training pairs, nearly perfect training results can be reached, i.e. the fidelity between the desired state and the network's output is nearly one. Further, we could observe that the training algorithm is robust to noise in the training data. Next, we presented an implementation of the DQNN training algorithm on a NISQ device, called DQNN<sub>NISQ</sub> [101], and tested it with the same learning task. Despite the high noise levels, we observed that the DQNN<sub>NISQ</sub> could generalise the information provided through the training data pairs.

Using the task of learning an unknown unitary, we compared the training success of DQNN<sub>NISQ</sub> with an implementation of the *quantum approximate optimisation algorithm* (QAOA) [103–105]. Whereas in the DQNN<sub>NISQ</sub> context, a perceptron acts on layers of different qubits, in the QAOA setting, a perceptron is defined as a sequence of operations, and all perceptrons act on the identical qubits. We could observe that both QNN architectures succeed in the learning task. However, the results indicate that the DQNN<sub>NISQ</sub> is more suitable for learning an unknown unitary operation compared to the QAOA.

As described in Section 3.5, many different models for a quantum perceptron and QNNs were proposed in the last years. An open research task is to perform a more exhaustive comparison, including more versions of QNNs and find out which QNN architectures best suits the training tasks. Further, in the following years, it will be essential to optimise the implementations of QNNs in general and the implementation of  $\text{DQNN}_{\text{NISQ}}$  in particular for quantum devices of coming generations.

### No free lunch theorem

To understand the ultimate limits for QNNs, we presented the *quantum no free lunch* (QNFL) theorem [100]. This result describes a bound on the probability that a quantum device, which can be modelled as a unitary process and is optimised with quantum data pairs, gives an incorrect output for a random input. The theorem enabled us to review the learning behaviour of QNNs.

For describing the proof of the QNFL theorem, we motivated a quantum risk by comparing the unknown unitary and the device approximated unitary using an average over all pure states. We compared the resulting bound with the classical *no free lunch* theorem for invertible functions. Further, we found that the results gained with the DQNN algorithm are close to achieving the QNFL bound. However, a slight discrepancy was expected since the process of evaluating the quantum risk includes empirical averages.

The here presented work on the QNFL theorem was already generalised by [298] to the case where these quantum states can be entangled to a reference system. Beyond that, the understanding of ultimate limits for quantum learning devices in different variations is key for further progress in the field of quantum machine learning.

### Quantum machine learning with graph-structured data

The above mentioned QNN architectures were trained with quantum data pairs. Moreover, we studied two extensions of this ansatz. The first one includes a possible underlying graph structure of the training data [101]. The structure of quantum devices leads to structured quantum data. Hence, we assumed to have knowledge about the underlying graph and found a training loss function including such a structure in the training process of the DQNNs and the  $\text{DQNN}_{\text{NISQ}}$ .

In the first two examples, we chose the desired output states to have a structure of a line graph and a graph describing a connected cluster. The input states were chosen to be random. In both examples, we could observe that including the graph structure in the training process leads to better generalisation results than a simple supervised ansatz.

Moreover, to also make the input states assigned to the graph structure, we used a synthetic classical graph, where each vertex was assigned to a label and an embedding vector. The graph was constructed via a classical deep walk. We used the embedding of a vertex to construct an input state and the labels for

## 8. Conclusion and outlook

constructing desired output states. Thus, both parts of a training data pair relied on the graph. Using a so constructed data set, we noticed that the training process leads to better results when including the knowledge on the graph structure, especially if only a few of the desired output states were supervised.

The in this work studied graphs were synthetically constructed. An exciting research direction would be to test the in [101] proposed graph-based loss function, using output data generated from actual quantum devices of which the build-in structure is known.

### Quantum generative adversarial networks

In the second extension, we discussed a different training goal. Whereas the original DQNN and the DQNN making use of graph structure were trained with data pairs in order to learn an underlying relation, we aimed for characterising a set of quantum states in order to extend it to quantum states having similar properties with the *discriminative quantum adversarial network* (DQGAN).

After describing the original classical *generative adversarial networks* (GANs), we explained how to build the DQGANs. Such an architecture is built of two DQNNs, where one of them plays a generative and one a discriminative role. The discriminator gets either the generator's output or a training data state as an input and is trained to distinguish well between these categories. The generator's goal is to trick the discriminator. These training goals can be realised in two different training losses, which are optimised successively in one training epoch.

Using the DQGAN, it is possible to extent a given data set to states with similar characteristics. We described a way to study the diversity of such a generated data set. In this regard, we found that, different to the training DQNN and DQNN with graph structure, more training epochs do not give better results. When training DQGANs, we could observe that a more extended training can lead to a maximum validation loss. However, after some training epochs, our examples showed that the generator only produces a small set of different states. We could assert that a carefully chosen number of training rounds can lead to both, a good validation loss and the generator's output diversity.

In this thesis, we only discussed a limited amount of examples for training data used to train DQGANs. The study of other data sets is of interest. One example could be a set of states with similar degrees of entanglement (concerning a chosen entanglement measure) [355]. Moreover, the application of the data sets produced by the generator is to study. In classical machine learning, the output of GANs is used, for example, for the training of other NN architectures, when training data is rare. Whether such an ansatz can be executed successfully using DQGANs, is left open as a future research topic. Further, GANs are not the only adaption of simple feed-forward NNs. It is also of interest to study the applications and behaviour of DQNNs in a recurrent or convolutional neural network settings. We leave these topics open for future work.

## Conclusion

This work introduced DQNNs, which are designed for fully quantum learning tasks, are capable of universal quantum computation and have low memory requirements while training. We showed via the QNFL theorem that these QNNs can be nearly optimally trained when learning an unknown unitary operation. Moreover, we demonstrated that DQNNs can make use of the graph structure of quantum data and can also be trained in a generative adversarial setting in order to extend quantum data sets.

Since we are still at the beginning of the age of quantum computers, this work leads to many interesting further research questions. We believe that QNNs will be essential for analysing and processing huge amounts of quantum data produced by the quantum devices of future generations. Today, we can already look back on many successes in the field of quantum computation in the last few years. Therefore we presume that quantum computing, including more qubits and less noise, compared to nowadays' NISQ devices, will be experimentally possible in the following decades. Quantum computers will not solve all of our world's problems. However, if applied wisely, we believe that they can lead to more exciting and fruitful applications in the coming time.



# Appendix





# A

---

## Dissipative quantum neural networks

---

In Section 4.5 the classical simulation of a *dissipative quantum neural network* (DQNN) was explained. The numerical simulation, including a generalisation analysis and noise robustness testing, was done with 2-3-2 DQNNs. In the following, we want to embrace the opportunity of an appendix and show some more examples. We study the following DQNN architectures: 1-1-1 in Figure A.1, 1-2-1 in Figure A.2, 2-3-2 in Figure A.3, 3-4-3 in Figure A.4, and 2-3-4-3-2 in Figure A.5. Note that the results in Figure A.3 were already discussed in Section 4.5. We replotted them here again for better comparison. All results are obtained with the code available at [276]. The training data pairs of all examples are of the form  $\{|\phi_x^{\text{in}}\rangle, |\phi_x^{\text{SV}}\rangle\}$  with  $|\phi_x^{\text{SV}}\rangle = Y |\phi_x^{\text{in}}\rangle$ , where  $Y$  is a unitary. In the plots values of the the training loss

$$\mathcal{L}_{\text{SV}} = \frac{1}{S} \sum_{x=1}^S F(|\phi_x^{\text{SV}}\rangle \langle \phi_x^{\text{SV}}|, \rho_x^{\text{out}}) = \frac{1}{S} \sum_{x=1}^S \langle \phi_x^{\text{SV}} | \rho_x^{\text{out}} | \phi_x^{\text{SV}} \rangle,$$

and the validation loss

$$\mathcal{L}_{\text{USV}} = \frac{1}{N-S} \sum_{x=S+1}^N \langle \phi_x^{\text{USV}} | \rho_x^{\text{out}} | \phi_x^{\text{USV}} \rangle,$$

with  $|\phi_x^{\text{USV}}\rangle = Y |\phi_x^{\text{in}}\rangle$  are depicted.

### Training

At first we comment on the evaluation of the loss function during a single training session. In all examples the validation and the training loss reach values near 1.

## Appendix

However, we can notice differences in the process: For training a 1-1-1 DQNN with  $S = 20$  training pairs and  $N - S = 80$  validation pairs the validation loss  $\mathcal{L}_{\text{USV}}$  reaches a value 0.95 at  $r_T = 382$ . For the other architectures the same value is reached at  $r_T = 164$  (1-2-1),  $r_T = 420$  (2-3-2), and  $r_T = 552$  (3-4-3). We can see as well that the distance between the validation and training loss is bigger.

But, especially in the first 500 rounds of training the 2-3-4-3-2, the validation loss is increasing rather slowly and a validation loss value of 0.95 is not reached until  $r_T = 795$ . Despite that, a fidelity of nearly 1 is reached at the end of the training. Note though that these values are not averaged over many training attempts and therefore are probably not very comparable, since every training process is unique due to the randomly assigned unitary initialisation and the generation of the training data pairs.

### Generalisation

A better comparison can be made for the different generalisation analysis plots. Here, every data point depicts an average over 10 independent training attempts.

For training a 1-1-1 DQNN in 1000 training epochs, the validation loss  $\mathcal{L}_{\text{USV}}$  reaches a value higher than 0.8 with only  $S = 2$  supervised training pairs. For the other architectures, the same value is reached at  $S = 2$  (1-2-1),  $S = 5$  (2-3-2),  $S = 9$  (3-4-3), and  $S = 6$  (2-3-4-3-2). As expected via the *quantum no free lunch* theorem, explained in Chapter 5, the generalisation of states of bigger dimensions general needs more training data pairs. Further, we want to point out that the values  $S = 5$  (2-3-2) and  $S = 6$  (2-3-4-3-2) lead to the assumption that additional layers do not improve the generalisation behaviour.

### Noise robustness

Analogously to the generalisation analysis, every data point in in the noise robustness plots is generated via an average over 10 independent training attempts.

Comparing the training of the 1-1-1 and the 1-2-1 DQNN, we can see that the additional qubit does not improve the resistance against noise. Further, we notice that the noise robustness shrinks with the number of qubits in the input layer. For  $\delta = 0.3$ , the 1-1-1 DQNN reaches a validation loss of 0.79 and the 1-2-1 a validation loss of 0.83. Using the same  $\delta$ , we reach validation loss values of 0.73 (2-3-2), 0.50 (2-3-4-3-2), and 0.62 (3-4-3) for the other DQNN architectures.









Figure A.1.: **1-1-1 DQNN**. The plots describe the training (a) and generalisation behaviour (b) as well as the noise robustness (c) of training the DQNN in  $k = 1000$  training epochs with  $\eta = 1$ ,  $\epsilon = 0.01$  and  $N = 100$  data pairs (based on a unitary  $Y \in \mathcal{U}(2)$ ).



(a) Training with  $S = 20$ .





(b) Generalisation analysis.



(c) Testing the noise robustness for  $S = 20$ .

Figure A.2.: **1-2-1 DQNN**. The plots describe the training (a) and generalisation behaviour (b) as well as the noise robustness (c) of training the DQNN in  $k = 1000$  training epochs with  $\eta = 1$ ,  $\epsilon = 0.01$  and  $N = 100$  data pairs (based on a unitary  $Y \in \mathcal{U}(4)$ ).

(a) Training with  $S = 20$ .

(b) Generalisation analysis.




(c) Testing the noise robustness for  $S = 20$ .


Figure A.3.: **2-3-2 DQNN**. The plots describe the training (a) and generalisation behaviour (b) as well as the noise robustness (c) of training the DQNN in  $k = 1000$  training epochs with  $\eta = 1$ ,  $\epsilon = 0.01$  and  $N = 100$  data pairs (based on a unitary  $Y \in \mathcal{U}(4)$ ).



(a) Training with  $S = 20$ .





(b) Generalisation analysis.



(c) Testing the noise robustness for  $S = 20$ .

Figure A.4.: **3-4-3 DQNN**. The plots describe the training (a) and generalisation behaviour (b) as well as the noise robustness (c) of training the DQNN in  $k = 1000$  training epochs with  $\eta = 1$ ,  $\epsilon = 0.01$  and  $N = 100$  data pairs (based on a unitary  $Y \in \mathcal{U}(9)$ ).

(a) Training with  $S = 20$ .

(b) Generalisation analysis.


(c) Testing the noise robustness for  $S = 20$ .

Figure A.5.: **2-3-4-3-2 DQNN**. The plots describe the training (a) and generalisation behaviour (b) as well as the noise robustness (c) of training the DQNN in  $k = 1000$  training epochs with  $\eta = 1$ ,  $\epsilon = 0.01$  and  $N = 100$  data pairs (based on a unitary  $Y \in \mathcal{U}(4)$ ).





# B

---

## Training with graph-structured quantum data

---

In Section 6.6, we described how to build a set of quantum data pairs with an underlying graph-structure using a classical walk. In the following, we plot the evolution of the validation loss after training a DQNN with

$$\mathcal{L}_{\text{SV+G}} = \mathcal{L}_{\text{SV}} + \gamma \mathcal{L}_G,$$

where

$$\mathcal{L}_{\text{SV}} \equiv \frac{1}{S} \sum_{x=1}^S \langle \phi_x^{\text{SV}} | \mathcal{E}(\rho_x^{\text{in}}) | \phi_x^{\text{SV}} \rangle,$$

and

$$\mathcal{L}_G \equiv \sum_{w,x \in V} [A]_{wx} d_{\text{HS}}(\mathcal{E}(\rho_w), \mathcal{E}(\rho_x)).$$

All results are generated with the code available at [276]. In Figure B.2,

$$\mathcal{L}_{\text{USV}} = \frac{1}{N-S} \sum_{x=S+1}^N \langle \phi_x^{\text{USV}} | \mathcal{E}(\rho_x^{\text{in}}) | \phi_x^{\text{USV}} \rangle$$

is plotted for a training with  $\gamma = 0$  and  $\gamma = -0.5$  using the graph-structured data built based on a classical deep walk. Since we discussed the generalisation behaviour of this training example only for  $1 \leq S \leq 10$  in Section 6.6, Figure B.2 plots the values of the loss functions after  $r_T = 1000$  training rounds for  $1 \leq S \leq 32$ . The figure shows that for cases with less than 15 of 32 supervised vertices, the validation loss is lower when ignoring the problem's graph structure. After that threshold both learning strategies are about equally good.

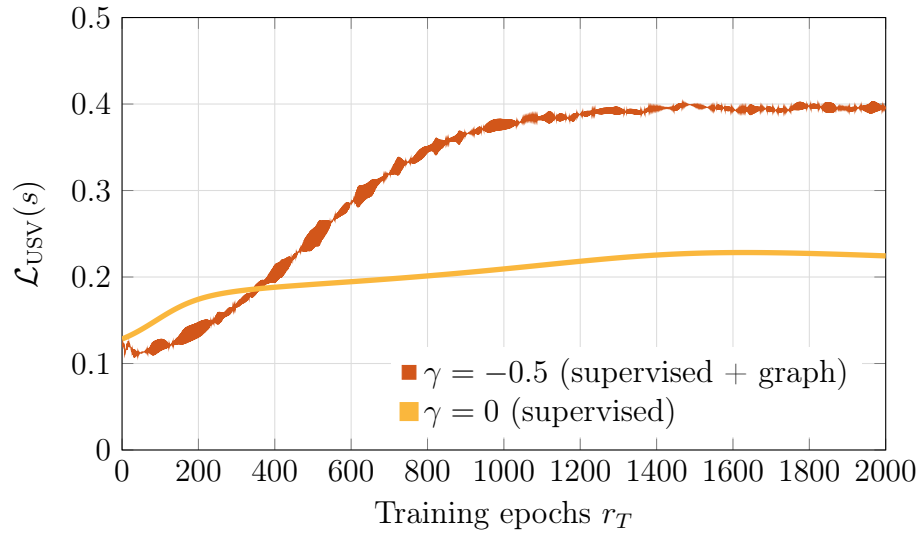
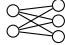


Figure B.1.: **Deep walk training.** The figure describes the training of a  DQNN ( $r_T = 2000$  epochs,  $\epsilon = 0.01$ ) trained with and without using the graph structure of a graph with 32 vertices produced by a classical deep walk.  $S = 10$  supervised data pairs are used.

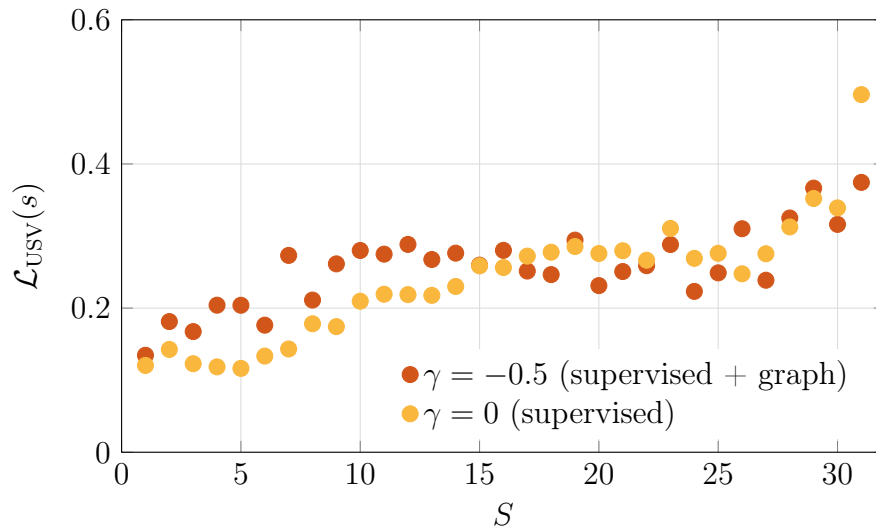
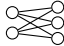


Figure B.2.: **Deep walk generalisation behaviour.** The plot describes the generalisation behaviour of a  DQNN ( $r_T = 1000$  epochs,  $\epsilon = 0.01$ ) trained with and without using the graph structure of a graph with 32 vertices produced by a classical deep walk. Each data point demonstrates an averaged over 10 independent training attempts.

# C

---

## Quantum generative adversarial networks

---

In Section 7.5 we discussed the classical simulation of the DQGAN algorithm. In the following we extend the numerical examples of this section. All results are generated with the code available at [276].


First, in Figure C.1 we compare the behaviour of the loss functions,

$$\begin{aligned}\mathcal{L}_D(\mathcal{E}_D, \mathcal{E}_G) &= \frac{1}{S} \sum_{x=1}^S \langle 0 | \mathcal{E}_D(\mathcal{E}_G(|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}|)) | 0 \rangle + \frac{1}{S} \sum_{x=1}^S \langle 1 | \mathcal{E}_D(|\phi_x^T\rangle \langle \phi_x^T|) | 1 \rangle, \\ \mathcal{L}_G(\mathcal{E}_D, \mathcal{E}_G) &= \frac{1}{S} \sum_{x=1}^S \langle 1 | \mathcal{E}_D(\mathcal{E}_G(|\psi_x^{\text{in}}\rangle \langle \psi_x^{\text{in}}|)) | 1 \rangle, \\ \mathcal{L}_V(\mathcal{E}_D, \mathcal{E}_G) &= \frac{1}{V} \sum_{i=v}^V \max_{x=1}^N (\langle \phi_x^T | \mathcal{E}_G(|\psi_i^{\text{in}}\rangle \langle \psi_i^{\text{in}}|) | \phi_x^T \rangle),\end{aligned}$$


for two different DQGAN architectures: 1-1-1 and 1-3-1. For the training we use  $S = 10$  states of the data sets

$$\begin{aligned}\text{data}_{\text{line}} &= \left\{ \frac{(N-x)|0\rangle + (x-1)|1\rangle}{\|(N-x)|0\rangle + (x-1)|1\rangle\|} \right\}_{x=1}^N, \\ \text{data}'_{\text{line}} &= \left\{ \frac{(N-x)|000\rangle + (x-1)|001\rangle}{\|(N-x)|000\rangle + (x-1)|001\rangle\|} \right\}_{x=1}^N,\end{aligned}$$

for  $N = 50$ , respectively. In both plots the opposed behaviour of  $\mathcal{L}_G$  and  $\mathcal{L}_D$  can be observed. Further, we notice that the validation loss does not reach values as high as in the case of the 1-1-1 network.




(a) Training a  $\circ-\circ-\circ$  DQGAN.




(b) Training a  $\circ-\circ-\circ$  DQGAN.

Figure C.1.: **Training a DQGAN.** The evolution of the training losses and validation loss during the training of a  $\circ-\circ-\circ$  (a) and a  $\circ-\circ-\circ$  (b) DQGAN in  $r_T = 1000$  epochs with  $\eta = 1$  and  $\epsilon = 0.01$  using 50 data pairs of the data sets  $\text{data}_{\text{line}}$  or  $\text{data}'_{\text{line}}$ , whereof 10 are used for training.



(a) Line trained with DQNN.



(b) Two clusters trained with DQNN.



(c) Two clusters plus  $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$  trained with DQNN.(d) Two clusters trained with DQNN<sub>Q</sub>.

Figure C.2.: **Diversity analysis of a DQGAN.** This plot describes the output's diversity of a  $\circ\text{-}\circ\text{-}\circ$  DQGAN (DQGAN<sub>Q</sub>) trained in 200 epochs with  $\eta = 1$  ( $\eta_D = 0.5, \eta_G = 0.1$ ) and  $\epsilon = 0.01$  ( $\epsilon = 0.25$ ) using 10 quantum states of the data sets  $\text{data}_{\text{line}}$  (a),  $\text{data}_{\text{cl}}$  (b,d) and  $\text{data}_{\text{cl}+}$  (c). The states from  $\text{data}_{\text{line}}$  are used as validation states.

So far, we have only described single runs of the DQGAN algorithm. In the following, we average in ten independent training attempts the data points of the

## Appendix

bar diagrams introduced in Section 7.5.

Figure C.2a depicts such an averaged bar diagram resulting after 200 training epochs of training a DQGAN with ten randomly chosen training states of the data set  $\text{data}_{\text{line}}$ . We can see that all elements of  $\text{data}_{\text{line}}$  get produced relatively equally.


Further, we extend our study with a second data set,

$$\text{data}_{\text{cl}} = \left\{ \frac{(2N-1)|0\rangle + (x-1)|1\rangle}{\|(2N-1)|0\rangle + (x-1)|1\rangle\|} \right\}_{x=1}^{\frac{N}{2}} \cup \left\{ \frac{(2N-1)|0\rangle + (x-1)|1\rangle}{\|(2N-1)|0\rangle + (x-1)|1\rangle\|} \right\}_{x=\frac{3N}{2}}^{2N}.$$


If we train a DQGAN and compare the generator's output with the  $\text{data}_{\text{cl}}$  data set, we expect that mainly some of the first and last states get produced. To study this, we train the DQGAN with  $S = 10$  randomly chosen training states of this set. Figure C.2b depicts the distribution of the generator's output after 200 training epochs. As expected, the generator does not produce all elements in  $\text{data}_{\text{line}}$  equally often. Due to the average of ten independent training attempts, the states  $|0\rangle$  and  $|1\rangle$  are very prominent in this plot. Since the state  $|0\rangle$  is produced more often, we assume that the training states randomly chosen in every training attempt the  $S = 10$  training data states were more often chosen of the first part of the cluster.

Further, by removing an arbitrary state of the data set  $\text{data}_{\text{cl}}$  and replacing it by  $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$  we obtain the connected cluster data set  $\text{data}_{\text{cl}+}$ . Figure C.2c shows the diversity of a generator resulting by training a DQGAN with this data set. We can see that some states in the middle are generated more often compared to the plot in Figure C.2b. However, the state  $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$  is not produced very often ( $x = 25$ ) and the resulting peak in the histogram is rather shifted more in the direction of the  $|1\rangle$  state ( $x = 50$ ).


In Figure C.2d the generator's diversity is depicted after a single training run of  $\text{DQGAN}_{\text{NISQ}}$  using the data set  $\text{data}_{\text{cl}}$  in  $r_T = 200$  training epochs and compared the output to the data set  $\text{data}_{\text{line}}$ . We can see that the generator is able to extend the clustered training data. However, similar to the training results of a DQNN shown in Figure C.2b, the  $\text{DQGAN}_{\text{NISQ}}$  does not achieve to produce the full range of training data. Note that we here use a slightly different implementation, compared to the one presented in Figure 4.11. The differences are marked with orange dashed lines in Figure C.3.




(a) Network.



(b) Implementation.



(c)  $U^1$  expressed by an unitary  $U_a^1$ , an additional unitary  $U_b^1$  and  $u$ -gates.



(d)  $U_a^1$  and  $U_b^1$  expressed by two-qubit unitaries.

Figure C.3.: **Implementation of a  $\text{DQNN}_{\text{NISQ}}^+$** . A  $\text{DQNN}_{\text{NISQ}}^+$  of the architecture 2-3-2<sup>+</sup> (a) can be implemented as quantum circuit using  $u$ -gates and unitary operations representing the layers of the network (b). Different to the  $\text{DQNN}_{\text{NISQ}}$ ,  $U^l$  is decomposed into two unitaries,  $U_a^1$  and  $U_b^1$ , and  $u$ -gates. The difference of a  $\text{DQNN}_{\text{NISQ}}^+$  compared to a  $\text{DQGAN}_{\text{NISQ}}$  is marked in orange dashed lines (c).  $U_a^1$  and  $U_b^1$  are expressed via two-qubit unitaries (d).





---

## Bibliography

---

- [1] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT press, 2012.
- [2] M. I. Jordan and T. M. Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015. doi:10.1126/science.aaa8415.
- [3] M. A. Nielsen. *Neural networks and deep learning*, volume 25. Determination press San Francisco, CA, 2015. URL <http://neuralnetworksanddeeplearning.com/>.
- [4] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. URL <http://www.deeplearningbook.org>.
- [5] J. Boyan, D. Freitag, and T. Joachims. A machine learning architecture for optimizing web search engines. In *AAAI Workshop on Internet Based Information Systems*, pages 1–8, 1996.
- [6] W. Serrano and E. Gelenbe. An intelligent internet search assistant based on the random neural network. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 141–153. Springer, 2016. doi:10.1007/978-3-319-44944-9\_13.
- [7] W. Serrano and E. Gelenbe. Intelligent search with deep learning clusters. In *2017 Intelligent Systems Conference (IntelliSys)*, pages 632–637, 2017. doi:10.1109/IntelliSys.2017.8324360.
- [8] W. Jiang and L. Zhang. Geospatial data to images: A deep-learning framework for traffic forecasting. *Tsinghua Science and Technology*, 24(1):52–64, 2019. doi:10.26599/TST.2018.9010033.
- [9] J. Ibarz and S. Banerjee. Updating google maps with deep learning and street view, 2017. URL <https://ai.googleblog.com/2017/05/Updating-google-maps-with-deep-learning.html>.
- [10] W. Awad and S. ELseuofi. Machine learning methods for spam e-mail classification. *International Journal of Computer Science & Information Technology*, 3(1):173–184, 2011.

## Bibliography

- [11] R. Hassanpour, E. Dogdu, R. Choupani, O. Goker, and N. Nazli. Phishing E-Mail Detection by Using Deep Learning Algorithms. In *Proceedings of the ACMSE 2018 Conference*, ACMSE '18, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3190645.3190719.
- [12] S. Bagui, D. Nandi, S. Bagui, and R. J. White. Classifying Phishing Email Using Machine Learning and Deep Learning. In *2019 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, pages 1–2, 2019. doi:10.1109/CyberSecPODS.2019.8885143.
- [13] E. G. Dada, J. S. Bassi, H. Chiroma, A. O. Adetunmbi, O. E. Ajibuwa, et al. Machine learning for email spam filtering: review, approaches and open research problems. *Heliyon*, 5(6):01802, 2019. doi:10.1016/j.heliyon.2019.e01802.
- [14] A. Kannan, K. Kurach, S. Ravi, T. Kaufmann, A. Tomkins, B. Miklos, G. Corrado, L. Lukacs, M. Ganea, P. Young, and V. Ramavajjala. Smart Reply: Automated Response Suggestion for Email. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 955–964, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2939672.2939801.
- [15] M. Henderson, R. Al-Rfou, B. Strobe, Y.-H. Sung, L. Lukács, R. Guo, S. Kumar, B. Miklos, and R. Kurzweil. Efficient natural language response suggestion for smart reply. 2017, arXiv:1705.00652.
- [16] Y. Weng, H. Zheng, F. Bell, and G. Tur. Occ: A smart reply system for efficient in-app communications. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2596–2603, 2019. doi:10.1145/3292500.3330694.
- [17] H. T. H. Nguyen, M. Wistuba, J. Grabocka, L. R. Drumond, and L. Schmidt-Thieme. Personalized Deep Learning for Tag Recommendation. In J. Kim, K. Shim, L. Cao, J.-G. Lee, X. Lin, and Y.-S. Moon, editors, *Advances in Knowledge Discovery and Data Mining*, pages 186–197, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-57454-7\_15.
- [18] Z. Qu, B. Li, X. Wang, S. Yin, and S. Zheng. An Efficient Recommendation Framework on Social Media Platforms Based on Deep Learning. In *2018 IEEE International Conference on Big Data and Smart Computing (BigComp)*, pages 599–602, 2018. doi:10.1109/BigComp.2018.00104.
- [19] S. Agrawal and A. Awekar. Deep Learning for Detecting Cyberbullying Across Multiple Social Media Platforms. In G. Pasi, B. Piwowarski, L. Azzopardi, and A. Hanbury, editors, *Advances in Information Retrieval*, pages 141–153, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-319-76941-7\_11.
- [20] J. Li, Q. Xu, N. Shah, and T. K. Mackey. A Machine Learning Approach for the Detection and Characterization of Illicit Drug Dealers on Instagram: Model Evaluation Study. 21(6):13803, 2021. doi:10.2196/13803.

- [21] E. D. Minin, C. Fink, H. Tenkanen, and T. Hiippala. Machine learning for tracking illegal wildlife trade on social media. *Nature Ecology and Evolution*, 2(3):406–407, 2018. doi:10.1038/s41559-018-0466-x.
- [22] K. R. Purba, D. Asirvatham, and R. K. Murugesan. Classification of instagram fake users using supervised machine learning algorithms. *International Journal of Electrical and Computer Engineering*, 10(3):2763, 2020. doi:10.11591/ijece.v10i3.pp2763-2772.
- [23] X. Liu and T. Zhu. Deep learning for constructing microblog behavior representation to identify social media users’ personality. *PeerJ Computer Science*, 2:e81, 2016. doi:10.7717/peerj-cs.81.
- [24] G. Gkotsis, A. Oellrich, S. Velupillai, M. Liakata, T. J. Hubbard, R. J. Dobson, and R. Dutta. Characterisation of mental health conditions in social media using Informed Deep Learning. *Scientific reports*, 7(1):1–11, 2017. doi:10.1038/srep45141.
- [25] D. Xue, L. Wu, Z. Hong, S. Guo, L. Gao, Z. Wu, X. Zhong, and J. Sun. Deep learning-based personality recognition from text posts of online social networks. *Applied Intelligence*, 48(11):4232–4246, 2018. doi:10.1007/s10489-018-1212-4.
- [26] R. Sawhney, P. Manchanda, P. Mathur, R. Shah, and R. Singh. Exploring and learning suicidal ideation connotations on social media with deep learning. In *Proceedings of the 9th workshop on computational approaches to subjectivity, sentiment and social media analysis*, pages 167–175, 2018. doi:10.18653/v1/W18-6223.
- [27] J. Kim, J. Lee, E. Park, and J. Han. A deep learning model for detecting mental illness from user content on social media. *Scientific Reports*, 10(1), 2020. doi:10.1038/s41598-020-68764-y.
- [28] S. Gao, V. D. Calhoun, and J. Sui. Machine learning in major depression: From classification to treatment outcome prediction. *CNS neuroscience & therapeutics*, 24(11):1037–1052, 2018. doi:10.1111/cns.13048.
- [29] F. Wang, L. P. Casalino, and D. Khullar. Deep learning in medicine - promise, progress, and challenges. *JAMA internal medicine*, 179(3):293–294, 2019. doi:10.1001/jamainternmed.2018.7117.
- [30] F. Piccialli, V. Di Somma, F. Giampaolo, S. Cuomo, and G. Fortino. A survey on deep learning in medicine: Why, how and when? *Information Fusion*, 66:111–137, 2021. doi:10.1016/j.inffus.2020.09.006.
- [31] D. Ravi, C. Wong, F. Deligianni, M. Berthelot, J. Andreu-Perez, B. Lo, and G.-Z. Yang. Deep Learning for Health Informatics. *IEEE Journal of Biomedical and Health Informatics*, 21(1):4–21, 2017. doi:10.1109/JBHI.2016.2636665.
- [32] M. Zitnik, M. Agrawal, and J. Leskovec. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 2018. doi:10.1093/bioinformatics/bty294.

## Bibliography

- [33] H. R. Roth, C. T. Lee, H.-C. Shin, A. Seff, L. Kim, J. Yao, L. Lu, and R. M. Summers. Anatomy-specific classification of medical images using deep convolutional nets. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pages 101–104, 2015. doi:[10.1109/ISBI.2015.7163826](https://doi.org/10.1109/ISBI.2015.7163826).
- [34] Z. Yan, Y. Zhan, Z. Peng, S. Liao, Y. Shinagawa, S. Zhang, D. N. Metaxas, and X. S. Zhou. Multi-Instance Deep Learning: Discover Discriminative Local Anatomies for Bodypart Recognition. *IEEE Transactions on Medical Imaging*, 35(5):1332–1343, 2016. doi:[10.1109/TMI.2016.2524985](https://doi.org/10.1109/TMI.2016.2524985).
- [35] M. Anthimopoulos, S. Christodoulidis, L. Ebner, A. Christe, and S. Mougiakakou. Lung pattern classification for interstitial lung diseases using a deep convolutional neural network. *IEEE transactions on medical imaging*, 35(5):1207–1216, 2016. doi:[10.1109/TMI.2016.2535865](https://doi.org/10.1109/TMI.2016.2535865).
- [36] Y. Cao, C. Liu, B. Liu, M. J. Brunette, N. Zhang, T. Sun, P. Zhang, J. Peinado, E. S. Garavito, L. L. Garcia, and W. H. Curioso. Improving Tuberculosis Diagnostics Using Deep Learning and Mobile Health Technologies among Resource-Poor and Marginalized Communities. In *2016 IEEE First International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*, pages 274–281, 2016. doi:[10.1109/CHASE.2016.18](https://doi.org/10.1109/CHASE.2016.18).
- [37] A. I. Khan, J. L. Shah, and M. M. Bhat. CoroNet: A deep neural network for detection and diagnosis of COVID-19 from chest x-ray images. *Computer Methods and Programs in Biomedicine*, 196:105581, 2020. doi:[10.1016/j.cmpb.2020.105581](https://doi.org/10.1016/j.cmpb.2020.105581).
- [38] G. Muhammad and M. Shamim Hossain. COVID-19 and Non-COVID-19 Classification using Multi-layers Fusion From Lung Ultrasound Images. *Information Fusion*, 72:80–88, 2021. doi:[10.1016/j.inffus.2021.02.013](https://doi.org/10.1016/j.inffus.2021.02.013).
- [39] W. Zhang, R. Li, H. Deng, L. Wang, W. Lin, S. Ji, and D. Shen. Deep convolutional neural networks for multi-modality iso-intense infant brain image segmentation. *NeuroImage*, 108:214–224, 2015. doi:[10.1016/j.neuroimage.2014.12.061](https://doi.org/10.1016/j.neuroimage.2014.12.061).
- [40] J. Kleesiek, G. Urban, A. Hubert, D. Schwarz, K. Maier-Hein, M. Bendszus, and A. Biller. Deep MRI brain extraction: A 3D convolutional neural network for skull stripping. *NeuroImage*, 129:460–469, 2016. doi:[10.1016/j.neuroimage.2016.01.024](https://doi.org/10.1016/j.neuroimage.2016.01.024).
- [41] Z. Ge, Z. Song, S. X. Ding, and B. Huang. Data mining and analytics in the process industry: The role of machine learning. *IEEE Access*, 5:20590–20616, 2017. doi:[10.1109/ACCESS.2017.2756872](https://doi.org/10.1109/ACCESS.2017.2756872).
- [42] T. D. Akinosho, L. O. Oyedele, M. Bilal, A. O. Ajayi, M. D. Delgado, O. O. Akinade, and A. A. Ahmed. Deep learning in the construction industry: A review of present status and future innovations. *Journal of Building Engineering*, page 101827, 2020. doi:[10.1016/j.jobe.2020.101827](https://doi.org/10.1016/j.jobe.2020.101827).
- [43] Q. Le, L. Miralles-Pechuán, S. Kulkarni, J. Su, and O. Boydell. An overview of deep learning in industry. *Data Analytics and AI*, pages 65–98, 2020.

- [44] P. A. O’Gorman and J. G. Dwyer. Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change, and extreme events. *Journal of Advances in Modeling Earth Systems*, 10(10):2548–2563, 2018. doi:10.1029/2018MS001351.
- [45] S. Ardabili, A. Mosavi, M. Dehghani, and A. R. Várkonyi-Kóczy. Deep learning and machine learning in hydrological processes climate change and earth systems a systematic review. In *International Conference on Global Research and Education*, pages 52–62. Springer, 2019. doi:10.1007/978-3-030-36841-8\_5.
- [46] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown, et al. Tackling climate change with machine learning. 2019, arXiv:1906.05433.
- [47] C. C. Aggarwal et al. *Neural networks and deep learning*, volume 10. Springer, 2018.
- [48] P. E. Ceruzzi, E. Paul, W. Aspray, et al. *A history of modern computing*. MIT press, 2003.
- [49] R. Schaller. Moore’s law: past, present and future. *IEEE Spectrum*, 34(6):52–59, 1997. doi:10.1109/6.591665.
- [50] E. Prati, D. Rotta, F. Sebastiano, and E. Charbon. From the Quantum Moore’s Law toward Silicon Based Universal Quantum Computing. In *2017 IEEE ICRC*, 2017. doi:10.1109/ICRC.2017.8123662.
- [51] M. Hilbert and P. López. The world’s technological capacity to store, communicate, and compute information. *Science*, 332(6025):60–65, 2011. doi:10.1126/science.1200970.
- [52] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [53] J. Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018. doi:10.22331/q-2018-08-06-79.
- [54] M. Brooks. Beyond quantum supremacy: the hunt for useful quantum computers. *Nature*, 574(7776):19–21, 2019. doi:10.1038/d41586-019-02936-3.
- [55] F. Arute, K. Arya, R. Babbush, D. Bacon, and et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. doi:10.1038/s41586-019-1666-5.
- [56] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 439(1907):553–558, 1992. doi:10.1098/rspa.1992.0167.
- [57] P. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994. doi:10.1109/SFCS.1994.365700.

## Bibliography

- [58] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996. doi:10.1145/237814.237866.
- [59] IBM Quantum, 2021. URL <https://quantum-computing.ibm.com>.
- [60] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017. doi:10.1038/nature23474.
- [61] V. Dunjko and H. J. Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports On Progress in Physics*, 81(7):074001, 2018. doi:10.1088/1361-6633/aab406.
- [62] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, et al. Variational quantum algorithms. *Nature Review Physics*, 3(9):625–644, 2020, arXiv:2012.09265. doi:10.1038/s42254-021-00348-9.
- [63] N. B. Lovett, C. Crosnier, M. Perarnau-Llobet, and B. C. Sanders. Differential Evolution for Many-Particle Adaptive Quantum Metrology. *Physical Review Letters*, 110(22):220501, 2013. doi:10.1103/PhysRevLett.110.220501.
- [64] M. Tiersch, E. Ganahl, and H. J. Briegel. Adaptive quantum computation in changing environments using projective simulation. *Scientific reports*, 5(1):1–18, 2015. doi:10.1038/srep12874.
- [65] G. Carleo and M. Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017. doi:10.1126/science.aag2302.
- [66] E. Aïmeur, G. Brassard, and S. Gambs. Quantum speed-up for unsupervised learning. *Machine Learning*, 90(2):261, 2013. doi:10.1007/s10994-012-5316-5.
- [67] G. D. Paparo, V. Dunjko, A. Makmal, M. A. Martin-Delgado, and H. J. Briegel. Quantum Speedup for Active Learning Agents. *Physical Review X*, 4(3):031002, 2014. doi:10.1103/PhysRevX.4.031002.
- [68] M. Schuld, I. Sinayskiy, and F. Petruccione. The quest for a Quantum Neural Network. *Quantum Inf. Process.*, 13(11):2567–2586, 2014. doi:10.1007/s11128-014-0809-8.
- [69] N. Wiebe, A. Kapoor, and K. M. Svore. Quantum Perceptron Models. 2016, arXiv:1602.04799.
- [70] M. Sasaki and A. Carlini. Quantum learning and universal quantum matching machine. *Physical Review A*, 66:022303, 2002. doi:10.1103/PhysRevA.66.022303.
- [71] S. Gambs. Quantum classification. *Physical Review Research* 2, 2(3):033125, 2008. doi:10.1103/PhysRevResearch.2.033125.
- [72] G. Sentís, J. Calsamiglia, R. Muñoz-Tapia, and E. Bagan. Quantum learning without quantum memory. *Scientific reports*, 2(1):1–8, 2012. doi:10.1038/srep00708.

- [73] V. Dunjko, J. M. Taylor, and H. J. Briegel. Quantum-Enhanced Machine Learning. *Physical Review Letters*, 117(13):130501, 2016. doi:10.1103/physrevlett.117.130501.
- [74] U. Alvarez-Rodriguez, L. Lamata, P. Escandell-Montero, J. D. Martín-Guerrero, and E. Solano. Supervised quantum learning without measurements. *Scientific reports*, 7(1):1–9, 2017. doi:10.1038/s41598-017-13378-0.
- [75] A. Monras, G. Sentís, and P. Wittek. Inductive supervised quantum learning. *Physical review letters*, 118(19):190503, 2017. doi:10.1103/PhysRevLett.118.190503.
- [76] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko. Quantum Boltzmann Machine. *Physical Review X*, 2018. doi:10.1103/PhysRevX.8.021050.
- [77] G. Sentís, A. Monràs, R. Muñoz Tapia, J. Calsamiglia, and E. Bagan. Unsupervised Classification of Quantum Data. *Physical Review X*, 2019. doi:10.1103/PhysRevX.9.041029.
- [78] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao. The Expressive Power of Parameterized Quantum Circuits. 2018, arXiv:1810.11922.
- [79] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf. Training deep quantum neural networks. *Nature Communications*, 11(1):1–6, 2020. doi:10.1038/s41467-020-14454-2.
- [80] M. Andrecut and M. Ali. A quantum neural network model. *International Journal of Modern Physics C*, 13(01):75–88, 2002. doi:10.1142/S0129183102002948.
- [81] W. R. d. Oliveira, A. J. Silva, T. B. Ludermir, A. Leonel, W. R. Galindo, and J. C. Pereira. Quantum Logical Neural Networks. In *10th Brazilian Symposium on Neural Networks*, pages 147–152, 2008. doi:10.1109/SBRN.2008.9.
- [82] M. Panella and G. Martinelli. Neural networks with quantum architecture and quantum learning. *International Journal of Circuit Theory and Applications*, 39(1):61–77, 2011. doi:10.1002/cta.619.
- [83] A. J. da Silva, T. B. Ludermir, and W. R. de Oliveira. Quantum perceptron over a field and neural network architecture selection in a quantum computer. 76:55–64, 2016. doi:10.1016/j.neunet.2016.01.002.
- [84] Y. Cao, G. G. Guerreschi, and A. Aspuru-Guzik. Quantum Neuron: an elementary building block for machine learning on quantum computers. 2017, arXiv:1711.11240.
- [85] K. H. Wan, O. Dahlsten, H. Kristjánsson, R. Gardner, and M. Kim. Quantum generalisation of feedforward neural networks. *npj Quantum Information*, 3(1):1–8, 2017. doi:10.1038/s41534-017-0032-4.
- [86] E. Farhi and H. Neven. Classification with Quantum Neural Networks on Near Term Processors. 2018, arXiv:1802.06002.
- [87] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd. Continuous-variable quantum neural networks. *Physical Review Research*, 2019. doi:10.1103/PhysRevResearch.1.033063.

## Bibliography

- [88] G. R. Steinbrecher, J. P. Olson, D. Englund, and J. Carolan. Quantum optical neural networks. *npj Quantum Information*, 5(1):1–9, 2019. doi:10.1038/s41534-019-0174-7.
- [89] E. Torrontegui and J. J. García-Ripoll. Unitary quantum perceptron as efficient universal approximator. *EPL (Europhysics Letters)*, 125(3):30004, 2019. doi:10.1209/0295-5075/125/30004.
- [90] F. Tacchino, P. Barkoutsos, C. Macchiavello, I. Tavernelli, D. Gerace, and D. Bajoni. Quantum implementation of an artificial feed-forward neural network. *Quantum Science and Technology*, 5(4):044010, 2020. doi:10.1088/2058-9565/abb8e4.
- [91] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib. Layerwise learning for quantum neural networks. *Quantum Machine Intelligence*, 3(1), 2020. doi:10.1007/s42484-020-00036-4.
- [92] K. Zhang, M.-H. Hsieh, L. Liu, and D. Tao. Toward trainability of quantum neural networks. 2020, arXiv:2011.06258.
- [93] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020. doi:10.1103/PhysRevA.101.032308.
- [94] K. Sharma, M. Cerezo, L. Cincio, and P. J. Coles. Trainability of Dissipative Perceptron-Based Quantum Neural Networks. 2020, arXiv:2005.12458.
- [95] Y. Zhang and Q. Ni. Design of quantum neuron model for quantum neural networks. *Quantum Engineering*, 3(3):e75, 2021. doi:10.1002/que2.75.
- [96] K. Vogel and H. Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Physical Review A*, 1989. doi:10.1103/PhysRevA.40.2847.
- [97] J. F. Poyatos, J. I. Cirac, and P. Zoller. Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate. *Physical Review Letters*, 1997. doi:10.1103/PhysRevLett.78.390.
- [98] M. Mohseni, A. T. Rezakhani, and D. A. Lidar. Quantum-process tomography: Resource analysis of different strategies. *Physical Review A*, 2008. doi:10.1103/PhysRevA.77.032322.
- [99] K. Beer, M. Khosla, J. Köhler, and T. J. Osborne. Quantum machine learning of graph-structured data. 2021, arXiv:2103.10837.
- [100] K. Poland, K. Beer, and T. J. Osborne. No free lunch for quantum machine learning, 2020, arXiv:2003.14103.
- [101] K. Beer, D. List, G. Müller, T. J. Osborne, and C. Struckmann. Training Quantum Neural Networks on NISQ Devices. 2021, arXiv:2104.06081.
- [102] K. Beer and G. Müller. Dissipative quantum generative adversarial networks. 2021, arXiv:2112.06088.



- [103] E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Approximate Optimization Algorithm. 2014, [arXiv:1411.4028](#).
- [104] E. Farhi and A. W. Harrow. Quantum Supremacy through the Quantum Approximate Optimization Algorithm. 2016, [arXiv:1602.07674](#).
- [105] S. Hadfield, Z. Wang, B. O’Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, 2019. [doi:10.3390/a12020034](#).
- [106] K. S. Saladin. *Human anatomy*. McGraw-Hill, 2011.
- [107] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature methods*, 10(5):413–420, 2013. [doi:10.1038/nmeth.2434](#).
- [108] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in neurons activity. *Bulletin of mathematical biophysics*, 5(115-133):10, 1943. [doi:10.1007/BF02478259](#).
- [109] S. C. Kleene. *Representation of events in nerve nets and finite automata*. Princeton University Press, 1956.
- [110] A. S. Householder and H. D. Landahl. *Mathematical biophysics of the central nervous system*. Number 1. Principia Press, Incorporated, 1958. [doi:10.1007/BF02478456](#).
- [111] B. Farley and W. Clark. Simulation of self-organizing systems by digital computer. *Transactions of the IRE Professional Group on Information Theory*, 4(4):76–84, 1954. [doi:10.1109/TIT.1954.1057468](#).
- [112] N. Rochester, J. Holland, L. Haibt, and W. Duda. Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Transactions on Information Theory*, 2(3):80–93, 1956. [doi:10.1109/TIT.1956.1056810](#).
- [113] F. Rosenblatt. Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Technical report, Cornell Aeronautical Lab Inc Buffalo NY, 1961.
- [114] W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller. Explaining Deep Neural Networks and Beyond: A Review of Methods and Applications. *Proceedings of the IEEE*, 109(3):247–278, 2021. [doi:10.1109/JPROC.2021.3060483](#).
- [115] F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958. [doi:10.1037/h0042519](#).
- [116] A. A. Minai and R. D. Williams. On the derivatives of the sigmoid. *Neural Networks*, 6(6):845–853, 1993. [doi:10.1016/S0893-6080\(05\)80129-7](#).
- [117] J. Mira and F. Sandoval. *From Natural to Artificial Neural Computation* , volume 930. Springer Science & Business Media, 1995.
- [118] G. Daqi and J. Yan. Classification methodologies of multilayer perceptrons with sigmoid activation functions. *Pattern Recognition*, 38(10):1469–1482, 2005. [doi:10.1016/j.patcog.2005.03.024](#).

## Bibliography

- [119] F. H. C. Tivive and A. Bouzerdoum. Texture Classification using Convolutional Neural Networks. In *TENCON 2006 - 2006 IEEE Region 10 Conference*, pages 1–4, 2006. doi:10.1109/TENCON.2006.343944.
- [120] M. L. Minsky and S. A. Papert. *Perceptrons*. MIT press, 1969.
- [121] R. Hecht-Nielsen. Neurocomputing: picking the human brain. *IEEE spectrum*, 25(3):36–41, 1988. doi:10.1109/6.4520.
- [122] I. Basheer and M. Hajmeer. Artificial neural networks: fundamentals, computing, design, and application. *Journal of Microbiological Methods*, 43(1):3–31, 2000. doi:10.1016/S0167-7012(00)00201-3. Neural Computing in Micrbiology.
- [123] S. Leijnen and F. v. Veen. The Neural Network Zoo. *Proceedings*, 47(1), 2020. doi:10.3390/proceedings2020047009.
- [124] G. Bebis and M. Georgiopoulos. Feed-forward neural networks. *IEEE Potentials*, 13(4):27–31, 1994. doi:10.1109/45.329294.
- [125] Z. Car, S. Baressi Šegota, N. Anđelić, I. Lorencin, and V. Mrzljak. Modeling the spread of COVID-19 infection using a multilayer perceptron. *Computational and mathematical methods in medicine*, 2020:1–10, 2020. doi:10.1155/2020/5714714.
- [126] K. Bhaskar and S. N. Singh. AWNN-Assisted Wind Power Forecasting Using Feed-Forward Neural Network. *IEEE Transactions on Sustainable Energy*, 3(2):306–315, 2012. doi:10.1109/TSTE.2011.2182215.
- [127] M. Omid, A. Baharlooei, and H. Ahmadi. Modeling Drying Kinetics of Pistachio Nuts with Multilayer Feed-Forward Neural Network. *Drying Technology*, 27(10):1069–1077, 2009. doi:10.1080/07373930903218602.
- [128] H. Ramchoun, M. A. J. Idrissi, Y. Ghanou, and M. Ettaouil. New modeling of multilayer perceptron architecture optimization with regularization: an application to pattern classification. *IAENG International Journal of Computer Science*, 44(3):261–269, 2017.
- [129] L. Jin, Z. Huang, Y. Li, Z. Sun, H. Li, and J. Zhang. On modified multi-output Chebyshev-polynomial feed-forward neural network for pattern classification of wine regions. *IEEE Access*, 7:1973–1980, 2018. doi:10.1109/ACCESS.2018.2885527.
- [130] P. Pełka and G. Dudek. Pattern-based forecasting monthly electricity demand using multilayer perceptron. In *International Conference on Artificial Intelligence and Soft Computing*, pages 663–672. Springer, 2019. doi:10.1007/978-3-030-20912-4\_60.
- [131] J. L. Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. doi:10.1207/s15516709cog1402\_1.
- [132] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur. Recurrent neural network based language model. In *Interspeech*, volume 2, pages 1045–1048. Makuhari, 2010.

- [133] M. Sundermeyer, R. Schlüter, and H. Ney. LSTM neural networks for language modeling. In *Thirteenth annual conference of the international speech communication association*, 2012.
- [134] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376, 2006. doi:10.1145/1143844.1143891.
- [135] A. Graves. Supervised sequence labelling. In *Supervised sequence labelling with recurrent neural networks*, pages 5–13. Springer, 2012. doi:10.1007/978-3-642-24797-2\_2.
- [136] A. Graves. Sequence transduction with recurrent neural networks. 2012, arXiv:1211.3711.
- [137] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt. Sequential deep learning for human action recognition. In *International workshop on human behavior understanding*, pages 29–39. Springer, 2011. doi:10.1007/978-3-642-25446-8\_4.
- [138] S. Ebrahimi Kahou, V. Michalski, K. Konda, R. Memisevic, and C. Pal. Recurrent neural networks for emotion recognition in video. In *Proceedings of the 2015 ACM on international conference on multimodal interaction*, pages 467–474, 2015. doi:10.1145/2818346.2830596.
- [139] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.
- [140] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. 2014, arXiv:1409.0473.
- [141] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [142] S. Lawrence, C. Giles, A. C. Tsoi, and A. Back. Face recognition: a convolutional neural-network approach. *IEEE Transactions on Neural Networks*, 8(1):98–113, 1997. doi:10.1109/72.554195.
- [143] N. Kalchbrenner, E. Grefenstette, and P. Blunsom. A convolutional neural network for modelling sentences. 2014, arXiv:1404.2188.
- [144] P. Kim. Convolutional neural network. In *MATLAB deep learning*, pages 121–147. Springer, 2017, arXiv:1404.2188.
- [145] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, 160(1):106–154, 1962. doi:10.1113/jphysiol.1962.sp006837.

## Bibliography

- [146] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi:10.1109/5.726791.
- [147] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [148] The MNIST database of handwritten digits. URL <http://yann.lecun.com/exdb/mnist/>.
- [149] E. L. Lehmann and G. Casella. *Theory of point estimation*. Springer Science & Business Media, 2006.
- [150] Z. Wang and A. C. Bovik. Mean squared error: Love it or leave it? A new look at Signal Fidelity Measures. *IEEE Signal Processing Magazine*, 26(1):98–117, 2009. doi:10.1109/MSP.2008.930649.
- [151] P. J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964. doi:10.1214/aoms/1177703732.
- [152] R. Rubinstein. The cross-entropy method for combinatorial and continuous optimization. *Methodology and computing in applied probability*, 1(2):127–190, 1999. doi:10.1023/A:1010091220143.
- [153] R. Y. Rubinstein. Combinatorial optimization, cross-entropy, ants and rare events. In *Stochastic optimization: algorithms and applications*, pages 303–363. Springer, 2001. doi:10.1007/978-1-4757-6594-6\_14.
- [154] R. Y. Rubinstein and D. P. Kroese. *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*, volume 133. Springer, 2004.
- [155] A. Tewari and P. L. Bartlett. On the Consistency of Multiclass Classification Methods. *Journal of Machine Learning Research*, 8(5), 2007. doi:10.1007/11503415\_10.
- [156] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality Reduction by Learning an Invariant Mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742, 2006. doi:10.1109/CVPR.2006.100.
- [157] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. URL [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2015/html/Schroff\\_FaceNet\\_A\\_Unified\\_2015\\_CVPR\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_2015/html/Schroff_FaceNet_A_Unified_2015_CVPR_paper.html).
- [158] M. F. Møller. *A scaled conjugate gradient algorithm for fast supervised learning*. Aarhus University, Computer Science Department, 1990. doi:10.1016/S0893-6080(05)80056-5.

- [159] M. I. Jordan and D. E. Rumelhart. Forward models: Supervised learning with a distal teacher. *Cognitive science*, 16(3):307–354, 1992. doi:10.1207/s15516709cog1402\_1.
- [160] R. Caruana and A. Niculescu-Mizil. An Empirical Comparison of Supervised Learning Algorithms. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, pages 161–168, New York, NY, USA, 2006. Association for Computing Machinery. doi:10.1145/1143844.1143865.
- [161] Q. V. Le. Building high-level features using large scale unsupervised learning. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8595–8598, 2013. doi:10.1109/ICASSP.2013.6639343.
- [162] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. 2015, arXiv:1511.06434.
- [163] N. Srivastava, E. Mansimov, and R. Salakhudinov. Unsupervised Learning of Video Representations using LSTMs. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 843–852, 2015. PMLR. URL <https://proceedings.mlr.press/v37/srivastava15.html>.
- [164] K. Wagstaff, C. Cardie, S. Rogers, S. Schrödl, et al. Constrained k-means clustering with background knowledge. In *Icml*, volume 1, pages 577–584, 2001.
- [165] K. Krishna and M. N. Murty. Genetic K-means algorithm. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 29(3):433–439, 1999. doi:10.1109/3477.764879.
- [166] A. K. Jain. Data clustering: 50 years beyond K-means. *Pattern recognition letters*, 31(8):651–666, 2010. doi:10.1016/j.patrec.2009.09.011.
- [167] P. Bholowalia and A. Kumar. EBK-means: A clustering technique based on elbow method and k-means in WSN. *International Journal of Computer Applications*, 105(9), 2014.
- [168] O. Chapelle, B. Scholkopf, and A. Zien, Eds. Semi-Supervised Learning (Chapelle, O. et al., Eds.; 2006) [Book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009. doi:10.1109/TNN.2009.2015974.
- [169] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589, 2014.
- [170] Y. Liu and K. Kirchhoff. Graph-based semi-supervised learning for phone and segment classification. In *INTERSPEECH*, pages 1840–1843, 2013.
- [171] R. Fergus, Y. Weiss, and A. Torralba. Semi-supervised learning in gigantic image collections. In *Advances in neural information processing systems*, pages 522–530, 2009.

## Bibliography

- [172] M. Shi and B. Zhang. Semi-supervised learning improves gene expression-based prediction of cancer recurrence. *Bioinformatics*, 27(21):3017–3023, 2011. doi:10.1093/bioinformatics/btr502.
- [173] X. Zhu, Z. Ghahramani, and J. D. Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *Proceedings of the 20th International conference on Machine learning (ICML-03)*, pages 912–919, 2003.
- [174] B. Liu, Z. Wu, H. Hu, and S. Lin. Deep metric transfer for label propagation with limited annotated data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. URL [https://openaccess.thecvf.com/content\\_ICCVW\\_2019/html/MDALC/Liu\\_Deep\\_Metric\\_Transfer\\_for\\_Label\\_Propagation\\_with\\_Limited\\_Annotated\\_Data\\_ICCVW\\_2019\\_paper.html](https://openaccess.thecvf.com/content_ICCVW_2019/html/MDALC/Liu_Deep_Metric_Transfer_for_Label_Propagation_with_Limited_Annotated_Data_ICCVW_2019_paper.html).
- [175] W. L. Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020. doi:10.2200/S01045ED1V01Y202009AIM046.
- [176] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [177] P. Werbos. *New tools for prediction and analysis in the behavioral sciences*. PhD thesis, 1974.
- [178] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. doi:10.1038/323533a0.
- [179] P. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. doi:10.1109/5.58337.
- [180] H. B. Curry. The method of steepest descent for non-linear minimization problems. *Quart. Appl. Math*, 2(3):258–261, 1944. doi:10.1090/qam/10667.
- [181] S. Du, J. Lee, H. Li, L. Wang, and X. Zhai. Gradient Descent Finds Global Minima of Deep Neural Networks. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1675–1685. 2019. URL <https://proceedings.mlr.press/v97/du19c.html>.
- [182] T. Schaul, I. Antonoglou, and D. Silver. Unit tests for stochastic optimization. 2013, arXiv:1312.6055.
- [183] D. Soydaner. A Comparison of Optimization Algorithms for Deep Learning. *International Journal of Pattern Recognition and Artificial Intelligence*, 34(13):2052013, 2020. doi:10.1142/S0218001420520138.
- [184] L. Bottou. Online algorithms and stochastic approximations. 1998. URL <http://leon.bottou.org/papers/bottou-98x>. revised, oct 2012.

- [185] M. Li, T. Zhang, Y. Chen, and A. J. Smola. Efficient Mini-Batch Training for Stochastic Optimization. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 661–670, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2623330.2623612.
- [186] B. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964. doi:10.1016/0041-5553(64)90137-5.
- [187] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In S. Dasgupta and D. McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1139–1147, 2013. PMLR. URL <https://proceedings.mlr.press/v28/sutskever13.html>.
- [188] R. A. Jacobs. Increased rates of convergence through learning rate adaptation. *Neural networks*, 1(4):295–307, 1988. doi:10.1016/0893-6080(88)90003-2.
- [189] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. 2014, arXiv:1412.6980.
- [190] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- [191] G. Hinton, N. Srivastava, and K. Swersky. Neural networks for machine learning lecture 6a overview of mini-batch gradient descent. *Cited on*, 14(8):2, 2012.
- [192] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 1982. doi:10.1007/BF02650179.
- [193] H. Hopf. Über die Abbildungen der dreidimensionalen Sphäre auf die Kugelfläche. In *Selecta Heinz Hopf*, pages 38–63. Springer, 1964. doi:10.1007/BF01457962.
- [194] D. Kaszlikowski, D. K. L. Oi, M. Christandl, K. Chang, A. Ekert, L. C. Kwek, and C. H. Oh. Quantum cryptography based on qutrit Bell inequalities. *Physical Review A*, 2003. doi:10.1103/PhysRevA.67.012310.
- [195] M. Kues, C. Reimer, P. Roztock, L. R. Cortés, S. Sciara, B. Wetzl, Y. Zhang, A. Cino, S. T. Chu, B. E. Little, et al. On-chip generation of high-dimensional entangled quantum states and their coherent control. *Nature*, 546(7660):622–626, 2017. doi:10.1038/nature22986.
- [196] M. V. Suslov, G. B. Lesovik, and G. Blatter. Quantum abacus for counting and factorizing numbers. *Physical Review A*, 2011. doi:10.1103/PhysRevA.83.052317.
- [197] A. R. Shlyakhov, V. V. Zemlyanov, M. V. Suslov, A. V. Lebedev, G. S. Paroanu, G. B. Lesovik, and G. Blatter. Quantum metrology with a transmon qutrit. *Physical Review A*, 2018. doi:10.1103/PhysRevA.97.022115.

## Bibliography

- [198] V. Soltamov, C. Kasper, A. Poshakinskiy, A. Anisimov, E. Mokhov, A. Sperlich, S. Tarasenko, P. Baranov, G. Astakhov, and V. Dyakonov. Excitation and coherent control of spin qudit modes in silicon carbide at room temperature. *Nature communications*, 10(1):1–8, 2019. doi:10.1038/s41467-019-09429-x.
- [199] S. L. Braunstein, A. Mann, and M. Revzen. Maximal violation of bell inequalities for mixed states. *Physical Review Letters*, 1992. doi:10.1103/PhysRevLett.68.3259.
- [200] A. Einstein, B. Podolsky, and N. Rosen. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review*, 1935. doi:10.1103/PhysRev.47.777.
- [201] A. Uhlmann. The "transition probability" in the state space of a\*-algebra. *Reports On Mathematical Physics*, 9(2):273–279, 1976. doi:10.1016/0034-4877(76)90060-4.
- [202] J. L. Dodd and M. A. Nielsen. A simple operational interpretation of the fidelity. 2001, arXiv:quant-ph/0111053.
- [203] M. Ozawa. Entanglement measures and the Hilbert-Schmidt distance. *Physics Letters A*, 268(3):158–160, 2000. doi:10.1016/S0375-9601(00)00171-7.
- [204] J. C. Garcia-Escartin and P. Chamorro-Posada. Swap test and Hong-Ou-Mandel effect are equivalent. *Physical Review A*, 87(5):052330, 2013. doi:10.1103/PhysRevA.87.052330.
- [205] L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles. Learning the quantum algorithm for state overlap. *New Journal of Physics*, 20(11):113022, 2018. doi:10.1088/1367-2630/aae94a.
- [206] C. M. Dawson and M. A. Nielsen. The Solovay-Kitaev algorithm. 2005, arXiv:quant-ph/0505030.
- [207] K. Kraus, A. Böhm, J. D. Dollard, and W. Wootters. States, effects, and operations: fundamental notions of quantum theory. Lectures in mathematical physics at the University of Texas at Austin. *Lecture notes in physics*, 190(190):IX–151, 1983.
- [208] J. Diestel and A. Spalsbury. *The joys of Haar measure*. American Mathematical Soc., 2014.
- [209] J. J. Duistermaat and J. A. Kolk. *Lie groups*. Springer Science & Business Media, 2012. doi:10.1007/978-3-642-56936-4.
- [210] A. W. Patrick Hayden, Debbie W. Leung. Aspects of generic entanglement. *Communications in mathematical physics*, 265(1):95–117, 2006. doi:10.1007/s00220-006-1535-6.
- [211] E. Pednault, J. Gunnels, D. Maslov, and J. Gambetta. On quantum supremacy, 2019. URL <https://www.ibm.com/blogs/research/2019/10/on-quantum-supremacy/>.
- [212] M. B. Hastings. A Classical Algorithm Which Also Beats  $\frac{1}{2} + \frac{2}{\pi} \frac{1}{\sqrt{D}}$  For High Girth MAX-CUT. 2021, arXiv:2111.12641.



- [213] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2063):3473–3482, 2005. doi:10.1098/rspa.2005.1546.
- [214] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342, 2011. doi:10.1145/1993636.1993682.
- [215] K. Fujii and T. Morimae. Commuting quantum circuits and complexity of Ising partition functions. *New Journal of Physics*, 2017. doi:10.1088/1367-2630/aa5fdb.
- [216] K. Fujii, H. Kobayashi, T. Morimae, H. Nishimura, S. Tamate, and S. Tani. Impossibility of Classically Simulating One-Clean-Qubit Model with Multiplicative Error. *Physical Review Letters*, 2018. doi:10.1103/PhysRevLett.120.200502.
- [217] R. Jozsa and M. V. den Nest. Classical simulation complexity of extended Clifford circuits. 2013, arXiv:1305.6190.
- [218] M. J. Bremner, A. Montanaro, and D. J. Shepherd. Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations. *Physical Review Letters*, 2016. doi:10.1103/PhysRevLett.117.080501.
- [219] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018. doi:10.1038/s41567-018-0124-x.
- [220] I. M. Georgescu, S. Ashhab, and F. Nori. Quantum simulation. *Reviews of Modern Physics*, 2014. doi:10.1103/RevModPhys.86.153.
- [221] D. Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985. doi:10.1098/rspa.1985.0070.
- [222] D. Coppersmith. An approximate Fourier transform useful in quantum factoring. 2002, arXiv:quant-ph/0201067.
- [223] J. P. Buhler, H. W. Lenstra, and C. Pomerance. Factoring integers with the number field sieve. *Lecture Notes in Mathematics*, 1554:50–94, 1993. doi:10.1007/BFb0091539.
- [224] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998. doi:10.1098/rspa.1998.0164.
- [225] M. B. Plenio and P. L. Knight. Decoherence limits to quantum computation using trapped ions. *Proceedings of the Royal Society A*, 453(1965):2017–2041, 1997. doi:10.1098/rspa.1997.0109.
- [226] M. McEwen, D. Kafri, Z. Chen, J. Atalaya, K. Satzinger, C. Quintana, P. V. Klimov, D. Sank, C. Gidney, A. Fowler, et al. Removing leakage-induced correlated errors in superconducting quantum error correction. *Nature communications*, 12(1):1–7, 2021. doi:10.1038/s41467-021-21982-y.

## Bibliography

- [227] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits. *Physical Review Letters*, 2016. doi:10.1103/PhysRevLett.117.060504.
- [228] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2):021314, 2019. doi:10.1063/1.5088164.
- [229] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014. doi:10.1038/nature13171.
- [230] A. D. Córcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow. Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nature communications*, 6(1):1–10, 2015. doi:10.1038/ncomms7979.
- [231] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. Girvin, L. Jiang, et al. Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature*, 536(7617):441–445, 2016. doi:10.1038/nature18949.
- [232] P. Shor. Fault-tolerant quantum computation. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 56–65, 1996. doi:10.1109/SFCS.1996.548464.
- [233] J. Preskill. Fault-tolerant quantum computation. In *Introduction to quantum computation and information*, pages 213–269. World Scientific, 1998. doi:10.1142/9789812385253\_0008.
- [234] J. Adcock, E. Allen, M. Day, S. Frick, J. Hinchliff, M. Johnson, S. Morley-Short, S. Pallister, A. Price, and S. Stanisic. Advances in quantum machine learning. 2015, arXiv:1512.02900.
- [235] E. Farhi, J. Goldstone, S. Gutmann, and H. Neven. Quantum Algorithms for Fixed Qubit Architectures. 2017, arXiv:1703.06199.
- [236] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig. Quantum machine learning: a classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2209):20170551, 2018. doi:10.1098/rspa.2017.0551.
- [237] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao. Expressive power of parametrized quantum circuits. *Physical Review Research*, 2(3):033125, 2020. doi:10.1103/PhysRevResearch.2.033125.
- [238] S. Kak. On quantum neural computing. *Information Sciences*, 83(3):143–160, 1995. doi:10.1016/0020-0255(94)00095-S.

- [239] T. Menneer and A. Narayanan. Quantum-inspired neural networks. *Tech. Rep. R329*, 1995.
- [240] M. Zak and C. P. Williams. Quantum neural nets. *International journal of theoretical physics*, 37(2):651–684, 1998. doi:10.1023/A:1026656110699.
- [241] M. Lewenstein. Quantum Perceptrons. *Journal of Modern Optics*, 41(12):2491–2501, 1994. doi:10.1080/09500349414552331.
- [242] M. Altaisky. Quantum neural network. 2001, arXiv:quant-ph/0107012.
- [243] M. Siomau. A quantum model for autonomous learning automata. *Quantum Information Processing*, 13(5):1211–1221, 2014. doi:10.1007/s11128-013-0723-5.
- [244] M. Schuld, I. Sinayskiy, and F. Petruccione. Simulating a perceptron on a quantum computer. *Physics Letters A*, 379(7):660–663, 2015. doi:10.1016/j.physleta.2014.11.061.
- [245] S. Mangini, F. Tacchino, D. Gerace, C. Macchiavello, and D. Bajoni. Quantum computing model of an artificial neuron with continuously valued input data. *Machine Learning: Science and Technology*, 2020. doi:10.1088/2632-2153/abaf98.
- [246] F. Tacchino, S. Mangini, P. K. Barkoutsos, C. Macchiavello, D. Gerace, I. Tavernelli, and D. Bajoni. Variational Learning for Quantum Artificial Neural Networks. *IEEE Transactions on Quantum Engineering*, 2:1–10, 2021. doi:10.1109/TQE.2021.3062494.
- [247] S. Mangini, F. Tacchino, D. Gerace, D. Bajoni, and C. Macchiavello. Quantum computing models for artificial neural networks. *EPL (Europhysics Letters)*, 2021. doi:10.1209/0295-5075/134/10002.
- [248] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Physical Review A*, 2018. doi:10.1103/PhysRevA.98.032309.
- [249] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019. doi:10.1088/2058-9565/ab5944.
- [250] K. Bu, D. E. Koh, L. Li, Q. Luo, and Y. Zhang. On the statistical complexity of quantum circuits. 2021, arXiv:2101.06154.
- [251] M. Ostaszewski, E. Grant, and M. Benedetti. Quantum circuit structure learning. 2019, arXiv:1905.09692.
- [252] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3):032331, 2019. doi:10.1103/PhysRevA.99.032331.
- [253] J. Stokes, J. Izaac, N. Killoran, and G. Carleo. Quantum Natural Gradient. *Quantum*, 2020. doi:10.22331/q-2020-05-25-269.
- [254] C. Xue, Z.-Y. Chen, Y.-C. Wu, and G.-P. Guo. Effects of quantum noise on quantum approximate optimization algorithm. 2019, arXiv:1909.02196.

## Bibliography

- [255] M. Alam, A. Ash-Saki, and S. Ghosh. Analysis of quantum approximate optimization algorithm under realistic noise in superconducting qubits. 2019, [arXiv:1907.09631](https://arxiv.org/abs/1907.09631).
- [256] F. B. Maciejewski, F. Baccari, Z. Zimborás, and M. Oszmaniec. Modeling and mitigation of realistic readout noise with applications to the Quantum Approximate Optimization Algorithm. *Quantum*, 5:464, 2021. doi:10.22331/q-2021-06-01-464.
- [257] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):1–6, 2018. doi:10.1038/s41467-018-07090-4.
- [258] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti. An initialization strategy for addressing barren plateaus in parametrized quantum circuits. *Quantum*, 3:214, 2019. doi:10.22331/q-2019-12-09-214.
- [259] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature communications*, 12(1), 2020. doi:10.1038/s41467-021-27045-6.
- [260] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles. Cost-function-dependent barren plateaus in shallow quantum neural networks. *Nature communication*, 12(1), 2020. doi:10.1038/s41467-021-21728-w.
- [261] G. Liu, W.-P. Ma, H. Cao, and L.-D. Lyu. A quantum Hopfield neural network model and image recognition. *Laser Physics Letters*, 2020. doi:10.1088/1612-202x/ab7347.
- [262] M. Schuld and N. Killoran. Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019. doi:10.1103/PhysRevLett.122.040504.
- [263] J. Romero, J. P. Olson, and A. Aspuru-Guzik. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, 2017. doi:https://doi.org/10.1088/2058-9565/aa8072.
- [264] A. Pepper, N. Tischler, and G. J. Pryde. Experimental realization of a quantum autoencoder: The compression of qutrits via machine learning. *Physical review letters*, 122(6):060501, 2019. doi:10.1103/PhysRevLett.122.060501.
- [265] D. Bondarenko and P. Feldmann. Quantum Autoencoders to Denoise Quantum Data. *Physical Review Letters*, 124:130502, 2020. doi:10.1103/PhysRevLett.124.130502.
- [266] T. Achache, L. Horesh, and J. Smolin. Denoising quantum states with Quantum Autoencoders—Theory and Applications. 2020, [arXiv:2012.14714](https://arxiv.org/abs/2012.14714).
- [267] G. Verdon, J. Pye, and M. Broughton. A Universal Training Algorithm for Quantum Deep Learning. 2018, [arXiv:1806.09729](https://arxiv.org/abs/1806.09729).
- [268] M. Sedlák, A. Bisio, and M. Ziman. Optimal Probabilistic Storage and Retrieval of Unitary Channels. *Physical Review Letters*, 2019. doi:10.1103/PhysRevLett.122.170502.

- [269] G. Verdon, T. McCourt, E. Luzhnica, V. Singh, S. Leichenauer, and J. Hidary. Quantum Graph Neural Networks. 2019, [arXiv:1909.12264](https://arxiv.org/abs/1909.12264).
- [270] I. Cong, S. Choi, and M. D. Lukin. Quantum convolutional neural networks. *Nature Physics*, 15(12):1273–1278, 2019. doi:10.1038/s41567-019-0648-8.
- [271] S. Arunachalam and R. de Wolf. A Survey of Quantum Learning Theory. 2017, [arXiv:1701.06806](https://arxiv.org/abs/1701.06806).
- [272] B. T. Kiani, S. Lloyd, and R. Maity. Learning unitaries by gradient descent. 2020, [arXiv:2001.11897](https://arxiv.org/abs/2001.11897).
- [273] M. R. Geller, Z. Holmes, P. J. Coles, and A. Sornborger. Experimental quantum learning of a spectral decomposition. *Phys. Rev. Research*, 2021. doi:10.1103/PhysRevResearch.3.033200.
- [274] M. M. Wolf. Mathematical foundations of supervised learning, 2018. URL [https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MA4801\\_2016S/ML\\_notes\\_main.pdf](https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MA4801_2016S/ML_notes_main.pdf).
- [275] J. Johansson, P. Nation, and F. Nori. QuTiP 2: A Python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 184(4):1234–1240, 2013. doi:10.1016/j.cpc.2012.11.019.
- [276] Quantum neural networks by Kerstin Beer - Python Code. URL [https://github.com/KerstinBeer/QuantumNeuralNetworks\\_Dissertation\\_PythonCode](https://github.com/KerstinBeer/QuantumNeuralNetworks_Dissertation_PythonCode).
- [277] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo. Theory of overparametrization in quantum neural networks. 2021, [arXiv:2109.11676](https://arxiv.org/abs/2109.11676).
- [278] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley. Geometric theory of nonlocal two-qubit operations. *Physical Review A*, 67(4):042313, 2003. doi:10.1103/PhysRevA.67.042313.
- [279] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley. Optimal quantum circuit synthesis from controlled-unitary gates. *Physical Review A*, 69(4):042309, 2004. doi:10.1103/PhysRevA.69.042309.
- [280] M. Blaauboer and R. De Visser. An analytical decomposition protocol for optimal implementation of two-qubit entangling gates. *Journal of Physics A: Mathematical and Theoretical*, 41(39):395307, 2008. doi:10.1088/1751-8113/41/39/395307.
- [281] P. Watts, M. O’Connor, and J. Vala. Metric structure of the space of two-qubit gates, perfect entanglers and quantum control. *Entropy*, 15(6):1963–1984, 2013. doi:10.3390/e15061963.
- [282] G. E. Crooks. Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition. 2019, [arXiv:1905.13311](https://arxiv.org/abs/1905.13311).
- [283] E. C. Peterson, G. E. Crooks, and R. S. Smith. Two-qubit circuit depth and the monodromy polytope. *Quantum*, 4:247, 2020. doi:10.22331/q-2020-03-26-247.

## Bibliography

- [284] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum Fingerprinting. *Physical Review Letters*, 87(16):167902, 2001. doi:10.1103/PhysRevLett.87.167902.
- [285] H. Abraham et al. Qiskit: An Open-source Framework for Quantum Computing. doi:10.5281/zenodo.2562110.
- [286] E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Approximate Optimization Algorithm Applied to a Bounded Occurrence Constraint Problem. 2015, arXiv:1412.6062.
- [287] D. Wecker, M. B. Hastings, and M. Troyer. Training a quantum optimizer. *Physical Review A*, 94(2):022309, 2016. doi:10.1103/PhysRevA.94.022309.
- [288] C. Y.-Y. Lin and Y. Zhu. Performance of QAOA on typical instances of constraint satisfaction problems with bounded degree. 2016, arXiv:1601.01744.
- [289] J. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, et al. Unsupervised machine learning on a hybrid quantum computer. 2017, arXiv:1712.05771.
- [290] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel. Quantum approximate optimization algorithm for maxcut: A fermionic view. *Physical Review A*, 97(2):022304, 2018. doi:10.1103/PhysRevA.97.022304.
- [291] M. Streif and M. Leib. Comparison of QAOA with quantum and simulated annealing. 2019, arXiv:1901.01903.
- [292] W. Lechner. Quantum Approximate Optimization With Parallelizable Gates. *IEEE Transactions on Quantum Engineering*, 1:1–6, 2020. doi:10.1109/TQE.2020.3034798.
- [293] S. Lloyd. Quantum approximate optimization is computationally universal. 2018, arXiv:1812.11075.
- [294] B. Nachman, M. Urbanek, W. A. de Jong, and C. W. Bauer. Unfolding quantum computer readout noise. *npj Quantum Information*, 6(1):1–7, 2020. doi:10.1038/s41534-020-00309-7.
- [295] E. Magesan, J. M. Gambetta, and J. Emerson. Characterizing quantum gates via randomized benchmarking. *Physical Review A*, 2012. doi:10.1103/PhysRevA.85.042311.
- [296] M. Sasaki, A. Carlini, and R. Jozsa. Quantum template matching. *Physical Review A*, 64:022317, 2001. doi:10.1103/PhysRevA.64.022317.
- [297] S. Gammelmark and K. Mølmer. Quantum learning by measurement and feedback. *New Journal of Physics*, 11(3):033017, 2009. doi:10.1088/1367-2630/11/3/033017.
- [298] K. Sharma, M. Cerezo, Z. Holmes, L. Cincio, A. Sornborger, and P. J. Coles. Reformulation of the no-free-lunch theorem for entangled data sets. 2020, arXiv:2007.04900.

- [299] D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1):67–82, 1997. doi:10.1109/4235.585893.
- [300] J. C. Bridgeman and C. T. Chubb. Hand-waving and interpretive dance: an introductory course on tensor networks. *Journal of physics A: Mathematical and theoretical*, 50(22):223001, 2017. doi:10.1088/1751-8121/aa6dc3.
- [301] I. Pitas. *Graph-based social media analysis*, volume 39. CRC Press, 2016.
- [302] T. S. Jepsen, C. S. Jensen, and T. D. Nielsen. Graph convolutional networks for road networks. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 460–463, 2019. doi:10.1145/3347146.3359094.
- [303] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020. doi:10.1016/j.aiopen.2021.01.001.
- [304] B. Perozzi, R. Al-Rfou, and S. Skiena. DeepWalk: Online Learning of Social Representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 701–710. Association for Computing Machinery, 2014. doi:10.1145/2623330.2623732.
- [305] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang. Network Embedding As Matrix Factorization: Unifying DeepWalk, LINE, PTE, and Node2Vec. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 459–467. Association for Computing Machinery, 2018. doi:10.1145/3159652.3159706.
- [306] S. Liu, M. F. Demirel, and Y. Liang. N-Gram Graph: Simple Unsupervised Representation for Graphs, with Applications to Molecules. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019, arXiv:1806.09206.
- [307] M. Khosla, J. Leonhardt, W. Nejdl, and A. Anand. Node representation learning for directed graphs. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 395–411. Springer, 2019. doi:10.1007/978-3-030-46150-8\_24.
- [308] M. Belkin, P. Niyogi, and V. Sindhvani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research*, 7(11):2399–2434, 2006.
- [309] T. N. Kipf, M. Welling, T. N. Kipf, and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*, 2017, arXiv:1609.02907.
- [310] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph Attention Networks. In *International Conference on Learning Representations*, 2018, arXiv:1710.10903.

## Bibliography

- [311] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs. 2017, [arXiv:1706.02216](#). [doi:1706.02216](#).
- [312] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*, pages 912–919, 2019, [arXiv:1810.00826](#).
- [313] S. Dernbach, A. Mohseni-Kabir, S. Pal, M. Gepner, and D. Towsley. Quantum walk neural networks with feature dependent coins. *Applied Network Science*, 4(1):76, 2019. [doi:10.1007/s41109-019-0188-2](#).
- [314] D. B. West et al. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River, 2001.
- [315] R. J. Trudeau. *Introduction to graph theory*. Courier Corporation, 2013.
- [316] B. Bollobás. *Modern graph theory*. Springer Science & Business Media, 2013.
- [317] M. Newman. *Networks*. Oxford university press, 2018. [doi:10.1093/oso/9780198805090.003.0012](#).
- [318] M. Khosla, V. Setty, and A. Anand. A Comparative Study for Unsupervised Network Representation Learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(5):1807–1818, 2021. [doi:10.1109/TKDE.2019.2951398](#).
- [319] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2021. [doi:10.1109/TNNLS.2020.2978386](#).
- [320] M. Rathee, Z. Zhang, T. Funke, M. Khosla, and A. Anand. Learnt Sparsification for Interpretable Graph Neural Networks. 2021, [arXiv:2106.12920](#).
- [321] S. Bhagat, G. Cormode, and S. Muthukrishnan. Node classification in social networks. In *Social network data analytics*, pages 115–148. Springer, 2011. [doi:10.1007/978-1-4419-8462-3\\_5](#).
- [322] Y. Rong, W. Huang, T. Xu, and J. Huang. Dropedge: Towards deep graph convolutional networks on node classification. 2019, [arXiv:1907.10903](#).
- [323] B. Li and D. Pi. Learning deep neural networks for node classification. *Expert Systems with Applications*, 137:324–334, 2019. [doi:10.1016/j.eswa.2019.07.006](#).
- [324] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. 2016, [arXiv:1609.02907](#).
- [325] P. Zhang, X. Zhou, P. Pelliccione, and H. Leung. RBF-MLMR: A Multi-Label Metamorphic Relation Prediction Approach Using RBF Neural Network. *IEEE Access*, 5:21791–21805, 2017. [doi:10.1109/ACCESS.2017.2758790](#).
- [326] K. Teru, E. Denis, and W. Hamilton. Inductive relation prediction by subgraph reasoning. In *International Conference on Machine Learning*, pages 9448–9457. PMLR, 2020. URL <http://proceedings.mlr.press/v119/teru20a.html>.



- [327] F. Liu, S. Xue, J. Wu, C. Zhou, W. Hu, C. Paris, S. Nepal, J. Yang, and P. S. Yu. Deep learning for community detection: progress, challenges and opportunities. 2020, [arXiv:2005.08225](#).
- [328] S. Pandit, D. H. Chau, S. Wang, and C. Faloutsos. Netprobe: a fast and scalable system for fraud detection in online auction networks. In *Proceedings of the 16th international conference on World Wide Web*, pages 201–210, 2007. [doi:10.1145/1242572.1242600](#).
- [329] F. Errica, M. Podda, D. Bacciu, and A. Micheli. A fair comparison of graph neural networks for graph classification. 2019, [arXiv:1912.09893](#).
- [330] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural Message Passing for Quantum Chemistry. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1263–1272. PMLR, 2017. URL <http://proceedings.mlr.press/v70/gilmer17a>.
- [331] X. Wang and A. Gupta. Videos as Space-Time Region Graphs. 2018.
- [332] S. Zhang, H. Tong, J. Xu, and R. Maciejewski. Graph convolutional networks: a comprehensive review. *Computational Social Networks*, 6(1):1–23, 2019. [doi:10.1186/s40649-019-0069-y](#).
- [333] S. Talaga and A. Nowak. Homophily as a process generating social networks: insights from social distance attachment model. 2019, [arXiv:1907.07055](#).
- [334] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. 2020, [arXiv:2006.11468](#).
- [335] C. Struckmann, Training Quantum Neural Networks with Graph-Structured Quantum Data. Master’s thesis, 2021.
- [336] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath. Generative Adversarial Networks: An Overview. *IEEE Signal Processing Magazine*, 35(1):53–65, 2018. [doi:10.1109/MSP.2017.2765202](#).
- [337] J.-Y. Zhu, P. Krähenbühl, E. Shechtman, and A. A. Efros. Generative visual manipulation on the natural image manifold. In *European conference on computer vision*, pages 597–613. Springer, 2016. [doi:10.1007/978-3-319-46454-1\\_36](#).
- [338] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved Techniques for Training GANs. *Advances in neural information processing systems*, 29:2234–2242, 2016.
- [339] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative Adversarial Text to Image Synthesis. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1060–1069, 2016. URL <http://proceedings.mlr.press/v48/reed16.html>.

## Bibliography

- [340] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. URL [https://openaccess.thecvf.com/content\\_cvpr\\_2017/html/Ledig\\_Photo-Realistic\\_Single\\_Image\\_CVPR\\_2017\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2017/html/Ledig_Photo-Realistic_Single_Image_CVPR_2017_paper.html).
- [341] P.-L. Dallaire-Demers and N. Killoran. Quantum generative adversarial networks. *Physical Review A*, 2018. doi:10.1103/PhysRevA.98.012324.
- [342] S. Lloyd and C. Weedbrook. Quantum Generative Adversarial Learning. *Physical Review Letters*, 2018. doi:10.1103/PhysRevLett.121.040502.
- [343] M. Benedetti, E. Grant, L. Wossnig, and S. Severini. Adversarial quantum circuit learning for pure state approximation. *New Journal of Physics*, 2019. doi:10.1088/1367-2630/ab14b5.
- [344] S. Chakrabarti, Y. Huang, T. Li, S. Feizi, and X. Wu. Quantum Wasserstein Generative Adversarial Networks. 2019, arXiv:1911.00111.
- [345] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou, and L. Sun. Quantum generative adversarial learning in a superconducting quantum circuit. *Science Advances*, 5(1):10, 2019. doi:10.1126/sciadv.aav2761.
- [346] C. Zoufal, A. Lucchi, and S. Woerner. Quantum Generative Adversarial Networks for learning and loading random distributions. *npj Quantum Information*, 5(1):103, 2019. doi:10.1038/s41534-019-0223-2.
- [347] C. Zoufal. *Generative Quantum Machine Learning*. PhD thesis, 2021, arXiv:2111.12738.
- [348] M. Y. Niu, A. Zlokapa, M. Broughton, S. Boixo, M. Mohseni, V. Smelyanskiy, and H. Neven. Entangling Quantum Generative Adversarial Networks. 2021, arXiv:2105.00080.
- [349] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017. URL <https://proceedings.mlr.press/v70/arjovsky17a.html>.
- [350] L. N. Vaserstein. Markov processes over denumerable products of spaces, describing large systems of automata. *Problemy Peredachi Informatsii*, 5(3):64–72, 1969.
- [351] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–2116, 2017. URL [https://openaccess.thecvf.com/content\\_cvpr\\_2017/html/Shrivastava\\_Learning\\_From\\_Simulated\\_CVPR\\_2017\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2017/html/Shrivastava_Learning_From_Simulated_CVPR_2017_paper.html).
- [352] M.-Y. Liu and O. Tuzel. Coupled generative adversarial networks. *Advances in neural information processing systems*, 29:469–477, 2016.

- [353] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. URL [https://openaccess.thecvf.com/content\\_cvpr\\_2017/html/Isola\\_Image-To-Image\\_Translation\\_With\\_CVPR\\_2017\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2017/html/Isola_Image-To-Image_Translation_With_CVPR_2017_paper.html).
- [354] G. Müller, Generative Adversarial Learning with Quantum Neural Networks. Master’s thesis, 2021.
- [355] L. Schatzki, A. Arrasmith, P. J. Coles, and M. Cerezo. Entangled Datasets for Quantum Machine Learning. 2021, [arXiv:2109.03400v2](https://arxiv.org/abs/2109.03400v2).



---

## Curriculum vitae

---

**Name:** Kerstin Beer  
**Date of birth:** 17.08.1993  
**Place of birth:** Hanover, Germany

### Academic career

- 2017 - 2021** **PhD student**, *Leibniz Universität Hannover*, Faculty of Mathematics and Physics, Institute for Theoretical Physics, Quantum information group, Prof. Dr. Tobias J. Osborne.  
Research, teaching, IT administration.
- 2015 - 2017** **Master of Science in Physics**, *Leibniz Universität Hannover*.  
Thesis: Contextuality and Cohomology
- 2012 - 2015** **Bachelor of Science in Physics**, *Leibniz Universität Hannover*.  
Thesis: Quantum Compiling - Zerlegung unitärer Quantenoperationen
- 2012 - 2017** **Scholarship**, *Heinrich Böll Stiftung*.
- 2004 - 2012** **Abitur**, *Albert-Einstein-Schule*, Laatzen.