Topology and localization out-of-equilibrium

DISS. ETH NO. 23762

TOPOLOGY AND LOCALIZATION OUT-OF-EQUILIBRIUM

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

EVERARD PIETER LODEWIJK VAN NIEUWENBURG

Master of Science, Universiteit Leiden, Leiden, Netherlands

born on 10.03.1990

citizen of
The Netherlands

accepted on the recommendation of

Prof. Dr. Sebastian D. Huber, examiner
Prof. Dr. Erez Berg, co-examiner

2016

The secret of getting ahead is getting started.

Mark Twain

SUMMARY

A recent trend in condensed matter physics is that of taking systems out-of-equilibrium. In such a setting, physical phenomena are possible that do not have equilibrium counterparts. Two typical scenarios for non-equilibrium physics are that of *open* systems, i.e. systems that couple to an environment, and that of quenches, where the system is suddenly taken out-of-equilibrium and its relaxation is studied. The present dissertation adresses three questions, two of which deal with these two scenarios of non-equilibrium situations. The third question is slightly orthogonal, and asks about the potential applications of machine learning to condensed matter physics.

As the first question, we asked about the fate of one dimensional symmetry protected topological (SPT) states when coupled to an environment. By showing how to extend the concept of entanglement spectra for pure states to that of mixed states, we were able to identify genuine mixed states with the properties of topologically non-trivial states. Additionally, we showed how to extend the classification scheme of SPT states based on projective symmetry representations to open systems. As an example, we have also considered the evolution of initial states that are topologically non-trivial upon coupling them to an environment. We have found that the topological non-trivial nature of the state dissapears on the way to reaching a steady state.

As the second question, we asked about the relaxation to equilibrium of a single spin probe coupled to an environment. We showed that the transient behaviour of the decoherence of this probe spin can be used to investigate an interacting and disordered spin chain. For weak disorder this spin chain is in a thermalizing phase, whereas for strong disorder the system avoids thermalization and instead forms a many-body localized (MBL) phase. We showed that the decoherence properties of the spin probe provide an unambiguous signal of the MBL phase. This work provides an important experimental probe for investigating the properties of many-body quantum systems further.

For the third question, we considered an equilibrium scenario and asked what potential use a machine learning algorithm can be for such cases. We showed that by compressing the information of the wavefunction in the form of the entanglement spectrum, it is possible to detect phase transitions using an approach based on artificial neural networks. The application of machine learning techniques to condensed matter physics is a very novel trend, and it is a fast-moving subject. Next to developing our own approach in this dissertation, we provide an overview of the recent works, thereby giving an understanding of the potential uses of machine learning in condensed matter physics.

These three projects together, each in their own right, address important questions in condensed matter physics. Each of them deals with a recent topic that is of interest to a broad community, and provides non-trivial results that further these communities.

KURZFASSUNG

Ein moderner Trend in der Festkörperphysik ist dass betrachten Nichtgleichgewichtssystemen. In einer solchen Situation gibt es physikalische Phänomene die in einem Gleichgewichtszustand nicht möglich sind. Zwei typische Fälle sind gegeben durch *offene* Systeme, das heisst Systeme die in Kontakt sind mit ihrer Umgebung, und Quenches, in denen das System plötzlich ausser Gleichgewicht gebracht wird und man die Relaxation studiert. In dieser Thesis werden drei Fragen beantwortet, wovon sich zwei mit nichtgleichtgewichts-systemen beschäftigen. Bei der dritten Frage handelt es sich um die Anwendung von Maschinellem Lernen auf die Festkörperphysik.

Als Erstes betrachten wir was mit eindimensionalen topologischen Zustände passiert wenn sie mit einer Umgebung in Kontakt gebracht werden. Dazu zeigen wir zuerst wie man das Konzept des Verschränkungsspekrums von reinen Zustände auf gemischte Zustände übertragen kann. Wir zeigen dass es gemischte Zustände gibt mit Verschränkngsspektra die topologisch nicht-triviale Situationen entsprechen. Ausserdem zeigen wir wie man die Klassifizierung solcher symmetriegeschützten topologischen Zustände, basierend auf projektiven Darstellungen, von reinen auf gemischte Zustände erweitern kann. Als Beispiel betrachten wir dazu die Zeitentwicklung eines Zustandes, nachdem dieser in einem reinen Zustand war und eine Wechselwirkung mit der Umgebung eingeschaltet wurde. Wir finden dass die nicht-triviale Natur des Zustandes eine solche Wechselwirkung mit der Umgebung nicht übersteht.

Die zweite Frage beschäftigt sich mit der Relaxation eines einzelnen Spins, nachdem dieser an einer Umgebung gekoppelt wurde. Wir zeigen dass man durch das Verhalten dieses Spins, nämlich durch seine Kohärenz, die Eigenschaften der Umgebung bestimmen kann. Wir verwenden diese Methode für eine wechselwirkende und ungeordnete Umgebung, die sich für schwache Unordnung als thermaler Zustand verhält, jedoch als vielteilchen-lokalisiert für starke Unordnung. Die Dekohärenz dieses einzelnen Spins setzt ein eindeutiges Signal der lokalisierten Phase voraus. Damit gibt dieses Projekt einen wichtigen Anhaltspunkt für experimentelle Untersuchungen in vielteilchen Quantensysteme.

Zuletzt untersuchen wir wie wir Maschinelles Lernen in der Festkörperphysik einsetzen können. Wir zeigen dass es möglich ist mittels eines neurales Netzwerks Phasenübergänge im entanglement spectra zu finden. Die Verwendung des Maschinellen Lernens ist ein neuer Trend welche sich momentan schnell weiterentwickelt. Unabhängig von unserer eigenen Arbeit geben wir einen Überblick über die Publikationen dieses jungen Feldes.

Diese drei Projekten behandeln zusammen, jedes in seinem eigenen Rahmen, wichtige Fragen der heutigen Festkörperphysik. Mit dieser Thesis hoffen wir relevante Antworten zu geben die ihren entsprechenden Forschungsgebieten einen Schritt weiter helfen.

PREFACE

It all began with Einstein's Special Relativity.

When the completion of my highschool studies drew closer and closer, the choice of what I wanted to study at University became more and more important. I don't make decisions very easily, but in the end decided to go for Computer Science. I quite enjoyed programming (and still do!), and liked to design clever algorithms. I choose to study at Leiden University, as opposed to the Technical University in Delft (mainly for the reason that programming in Leiden was taught in `C++` while Delft had adopted `Java`...). This choice turned out to be a very important one. In the last few months at highschool, I happened to be introduced to the subject of special relativity. It was that moment at which I decided I would regret not also trying out Physics. I was fascinated by the counterintuitive facts about the world that special relativity imposed on it, and so I decided to also enroll in Physics. What follows was a year packed with overlapping courses and schedules, where I tried to absorb as much of both fields as possible. After this first year though, I was hooked mostly to Physics. I successfully completed the first year curriculum of both fields, but decided to focus my attention on Physics from then on. I did continue taking courses in Computer Science, but choose only those topics that really interested me, such as Artificial Intelligence and Machine Learning.

Today, I am happy for the choice I made back in highschool, and certainly do not regret it. I get to enjoy very interesting research in Physics, and can continue to enjoy programming in the form of numerical simulations of physical systems. I even get to pursue the interesting combination of physics and machine learning.

I hope I am able to convey this sense of excitement to you in the rest of this thesis.

Everard Pieter Lodewijk van Nieuwenburg
Zürich, July 2016

CONTENTS

I

CONDENSED MATTER THEORY

1

General Introduction

Condensed matter theory is a broad field, encompassing many different topics that each would deserve a separate introduction. The aim of this chapter is not to provide allround introductions to all of such topics, but rather to put the relevant concepts of this thesis into context. These topics include topological states and symmetry breaking phases of matter, open systems and non-equilibrium quantum physics, and entanglement and matrix product states. In addition, we shortly touch upon the topic of machine learning applications in condensed matter physics.

1

1.1. THESIS LAYOUT

This thesis is divided into several *parts*, each covering one or multiple chapters that share the main concept indicated by the part title. It is the purpose of this section to provide a short summary of these logical divisions. Hopefully, this may serve as a guide to how the parts are related.

- Part I: Condensed Matter Theory
  A general overview of the concepts is provided in this part, setting the stage and providing the background for the rest of this thesis. The following parts and chapters assume a certain broad overview of such relevant topics, and the aim of this part is to provide exactly that. We mainly introduce the general field of condensed matter physics, and highlights topics such as topology, entanglement and open systems. Additionaly, we provide a quick outlook on recent developments in the use of machine learning for analyzing condensed matter physics data.

- Part II: Equilibrium
  The emphasis in part II is on analysis of states in equilibrium. We introduce the concept of symmetry protected topological states and their classification based on projective symmetries. A useful way of doing so is by making use of the matrix product state formalism, which is also partially discussed in this part. A rigorous introduction to matrix product states can be found in Appendix A. The projective symmetry analysis will be extended upon in part III, and Appendix B has been devoted to provide more detail on this. Also part of the equilibrium analysis is the use of machine learning to identify phases of matter. We give an overview of general uses of machine learning in (condensed matter) physics, and develop a new method for detecting phase transitions.

- Part III: Out-of-Equilibrium
  Part III deals with analysis of out-of-equilibrium systems. In particular, we consider the effects of an external bath on topological states, and propose an extension of the projective symmetry analysis to general mixed states. Additionally, we introduce many-body localization initially as an equilibrium case, but argue that regular equilibrium thermodynamics does not detect it. Hence we study the out-of-equilibrium transient behaviour of a probe spin coupled to a many-body localized environment, and show that it is capable of detection.

- Part IV: Conclusion and Appendices
  In the conclusion, we discuss the obtained results and conclusions in a critical manner. The appendices are not only meant to provide rigorous derivations, but instead to provide alternative viewpoints on certain approaches. Especially in Appendix A, we provide an alternate introduction to explain matrix product states based on finite state machines. These sections have been delegated to an appendix in order for them not to obstruct the flow of the story.

Part I is in principle logically independent of the other sections, and will continue to

introduce the concepts and their bigger picture. Part II and III can be read as stand-alone without referring to part I, but the results in part III make use of the ideas described in detail in part II. An important remark is that the Appendices are not meant as 'the place with lengthy derivations', but rather should provide different viewpoints that are complementary to the main text but would distract from the main story. As an example, I think the finite-state-machine viewpoint on matrix product states is very powerful and especially insightful, but it is by no means necessary to understand results in parts II and III. The conclusions in part IV are meant to be a critical look at the obtained results, and suggests improvements and continuations on the works described in general.

1.2. CONDENSED MATTER PHYSICS

Since this dissertation deals with topics from condensed matter physics, let us sketch the broader picture of how this field is related to other fields of physics in general. The reader may wish to skip to sections 1.3 and 1.4 for topics more specific to the present thesis.

In any case, we quickly introduce these concepts and a little history here to fall back on in the beginning of Chapter 2. In that chapter, we will explicitly discuss a particular quantum state called the 'AKLT'-state, and use it to introduce in more detail the topics touched upon here. Especially the symmetry breaking and non-symmetry breaking (topological) phases of matter and the concept of entanglement will play an essential role in the rest of this dissertation. In Chapter 3, we will again build upon these concepts even further, and ask if we can teach them to a machine.

First, however, to put the entirety of this work into context, we continue with a review of condensed matter physics in its broadest sense. A definition that covers the most imporant content could read:

> *Quantum condensed matter physics is the subdiscipline of physics focussing on the properties of macroscopic objects in terms of the collective behaviours of (near) infinities of its microscopic, quantum-mechanical constituents.*

Of course, there are cases that fall outside such a broad definition. Subjects that stay on the level of microscopic constituents, dealing with single- or few-body physics are often classified as condensed matter physics even though they do not focus on macroscopic objects. In most of those cases however, the single particles such as electrons are in the presence of a potential that comes from a macroscopic number of atoms, justifying their classification as condensed matter physics. This should be contrasted to the study of free, high-energy particles and their properties, which rather belongs to the realm of particle physics.

One may wonder why this definition for condensed matter physics focuses on the study of macroscopic objects. After all, media coverage in the current age often highlights

progress in particle physics and its aims to try and break things up into their smallest possible building blocks. Once we understand those, is it not true we may deduce all else from that?

The answer to this question, for all practical purposes, is a definitive no. The behaviour of a collection of many particles is *qualitatively* different from that of its constituents [1]. The simplest and perhaps most abundant example that comes to mind is the difference between liquid water and ice: both are made from water molecules, yet one is rigid while the other is not. These new properties of collections of particles are referred to as *emergent* properties, arising solely due to the sheer number of constituents and their interactions. The constituents start to *order* according to a specific pattern, and such an ordered state is generally called a phase of matter. The example of ice given before is a phase of matter with spatial ordering, but so too are (anti-)ferromagnets in which magnetic moments (anti-)align, or even superconductors in which electrons order themselves in pairs.

1.2.1. RENORMALIZATION FLOW

The typical line of thought in getting to such collective and emergent behaviour is by starting from 'high' energy physics, sometimes also called the ultraviolet sector of a physical theory. For condensed matter, this starting point is usually that of individual electrons and atoms. They are described by operators for position, momentum and possibly include a spin degree of freedom. When many atoms and electrons are brought together, collective behaviour starts to emerge on energy scales that are much lower than that of single atoms. The atoms start forming a crystal, and low energy modes describing lattice distortions start becoming available on a spatial scale that is much larger than that of the single atoms [2]. It turns out that this decrease of energy scales and increase in spatial scales is a universal behaviour for emergent phenomena. Indeed, one may formulate a well-known physical formalism based on this principle: renormalization flow.

Along this flow to lower and lower energy scales, a few different scenarios may occur. An evident possibility perhaps, is that the system stops changing when we consider another increase in spatial scale or decrease in energy scale. The system is then self-similar, and is said to be in a critical state. In this state, by definition of self-similarity, the system has no unique length scale nor energy scale. Contrary to what one may initially believe, the critical state is actually very common: classical phase transitions between phases of matter are described by such critical states [3].

Apart from critical states, it may happen that during this flow process there is the formation of 'bound states'. The name bound state generally refers to collective objects with well defined properties. As an example we may think of the lattice distortions that were mentioned before, which form when we consider a multitude of atoms forming a crystal. A new effictive theory describing just these *phononic* modes may be set up without regard to the underlying atoms [4]. An important property of bound states however, is that

they may cost a finite amount of energy to form and hence to break up [1]. That means that as long as we consider energy scales that are smaller than this energy cost, called the energy gap or mass gap, we may formulate a new effective theory for the collective objects and their mutual interactions. Another typical and illustrative example of this (albeit less relevant for condensed matter) is that of atoms themselves. Atomic physics is an effective theory of bound nuclei [5] (which in turn are bound states of protons and neutrons), and one no longer has to deal with those nucleic constituents directly as long as energy scales are considered that are less than needed to break the atom apart.

1.2.2. ADIABATIC CONTINUITY

Another very useful idea in condensed matter physics is that of adiabatic continuity [6]. In essence, this means that it really does not matter very much what exactly we choose as the ultraviolet starting point. As long as we include the important features, such as symmetries, conservation laws, etc., the effective theories we end up with are qualitatively identical. Quantitatively they may well differ, but the claim is that these differences in numbers can be resolved adiabatically: we can tweak parameters to get to the correct quantitative result without making qualitatitve changes. Moreover, we should realize that especially in condensed matter theory it is rarely actually possible to get the numbers correct a priori. Most likely, these numbers are substituted by general variables and later taken from experiments.

That means that in almost every case, the starting point of calculations will be a simplified Hamiltonian description of the system that captures the essence of the microscopic model. It is in principle not difficult to write down a Hamiltonian model for the full set of particles and their interactions (although it is in practice impossible to solve it), rather it means that we can arrive at qualitatively correct results by approximating it to something that is much simpler.

1.2.3. SYMMETRY BREAKING

Such simplified models must at the very least preserve the symmetries of the original model, bringing us to one of the most important principles in physics: the spontaneous symmetry breaking paradigm by Ginzburg and Landau [7]. On a quantum mechanical level, nature is dominated by symmetry. In daily life this concept is not very apparent, since we are not used to thinking about liquid water having more symmetry than ice. Whereas quantum mechanical eigenstates can be ordered according to symmetries they possess (or don't possess), for classical objects this dominating role of symmetries is absent. The script that first first introduced me to condensed matter physics [8], written by Prof. Jan Zaanen, says this quite directly: *Classical objects do not correspond to quantum-mechanical eigenstates.*

[1] In the example of phonons there are those types of phononic modes in crystals for which long wavelength modes cost no energy, and those that do cost energy to excite. I refer to the latter of the two here.

The concept of symmetry breaking is ubiquitous in nature. Let us get back to the example of liquid water versus ice. It is fairly straightforward to write down a quantum mechanical model for the water molecules. This model will have translational symmetry, meaning we can displace a molecule to any position and the model will look the same. The eigenstates of the Hamiltonian describing this system therefore also obey this invariance under translations, meaning it is impossible to determine the locations of the atoms from the eigenstates themselves. This is not the case for the everyday ice crystals, in which the molecules order according to a fixed pattern and we are able to directly point at them. The translational invariance on the level of the single constituents got lost (is 'broken'), and instead we are left with translational invariance only over discrete distances, and of the *entire* crystal as a whole. This crystalline state is hence said to be a symmetry broken state, having broken full translational invariance to only a discrete version.

This phenomenon of symmetry breaking is a fundamental principle in physics. States of matter can be classified according to the sets of symmetries they possess and/or break. The textbook example of a ferromagnet demonstrates this clearly. We consider magnetic moments that can (anti-)align themselves along an axis that we will call the z-axis. These magnetic moments are due to the spin degree of freedom of particles, and we refer to their directions as spin-'up' and spin-'down'. The Hamiltonian describing this system is invariant under 180° flips of all the spins. After all, the Hamiltonian only acts between neighbouring spins, and by flipping both of them we will not be able to change their relative alignment. Hence eigenstates have this symmetry, too. It is clear however, that the lowest energy eigenstate of this system is one in which all these moments have aligned. In fact, in the absence of an external magnetic field that chooses a direction, that means there are two possible groundstates: one in which all the spin align in the spin-down direction and one in which they align in the spin-up direction. The quantum mechanical system *chooses* one of these states spontaneously (dominated by random fluctuations in the environment). But flipping all of the magnetic moments transforms one groundstate into the other, and not back into itself as invariance would dictate. This groundstate is therefore said to spontaneously break this symmetry. In other words, the ferromagnet is characterized by a broken spin-rotational symmetry [2].

Instead of analyzing states solely by the breaking of symmetries, we may additionally define *measureable* quantities that detect this spontaneous symmetry breaking. Such quantities are refered to as order parameters, and are chosen such that a non-zero expectation value of the order parameter indicates the presence of the corresponding order. For the ferromagnetic case the order parameter is simply the total magnetization of the state, whereas for the liquid to solid transition in water the usual order parameter is the density difference between the two phases [3].

[2] This is a special case, since the ferromagnetic state is actually an eigenstate of the Hamiltonian. For an antiferromagnet, the classical state is characterized by alternating spin-up and spin-down. This state is however not an eigenstate of the Hamiltonian that favors anti-alignment of the magnetic moments.

1.2.4. TOPOLOGY

Until not so long ago, it was thought that all phases of matter could be classified according to this paradigm of symmetry breaking. Over the past few decades however, the focus has shifted towards phases that can not be classified according to symmetry breaking [9, 10]. There are phases of matter that cannot be distinguished from each other by symmetries, but which have well-defined differences and are separated by energy gap closing transitions. These states are called *topologically* distinct, and we will discuss the one dimensional variants in more detail in a following chapter. The type of topological phases that can exist in 1D are more restricted than general topological phases. In 1D topological phases only exist in the presence of symmetries, and are hence called symmetry protected topological (SPT) phases [11].

1.2.5. ENTANGLEMENT

A concept that has gained a lot of attention over the last few years in condensed matter physics (but originating in the field of quantum information theory), is that of entanglement [12]. Entanglement has especially been a useful concept in the analysis of many-body quantum physics, and has led to both theoretical as well as significant numerical progress. The essence of entanglement is that two entangled objects cannot be treated as separate, even if they find themselves on opposite edges of the universe. A measurement performed on one of the two instantaneously affects the state of the other.

Wavefunctions describing objects that are not *too* entangled (a statement to be made more precise in Chapter 2 and Appendix A), can be efficiently described using the formalism of matrix product states (MPS). Entanglement has turned out to be a key feature distinguishing different 1D topological phases and in the understanding of localizing phases (which will be discussed in Chapter 5).

1.2.6. LOCALIZING PHASES

If one prepares a classical system in an ordered state and leaves it by itself, thermodynamics will wash out this initial order and thermalize the system with its environment [13]. Thermalization means that measurement outcomes on the system can be described by assigning it to a Gibbs ensemble, i.e. the probability of the system being in a given configuration is weighted by a Boltzmann factor involving the energy of the respective configuration. Thermalization is hard to prevent, although preventing it may be desirable if one wishes to preserve the information contained in an initial state of a system [14]. Of course, it is in principle possible to construct a system that has a number of conserved quantities in which one may store information. This most likely involves an extreme fine-tuning of the system's parameters, making a practical realisation of such an *integrable* system impossible [15].

Quantum systems have another way of preventing thermalization however, that is in a

sense more robust than hard-coded integrability. Quantum mechanics namely seems to provide a way of preventing thermalization due to the presence of disorder and interactions. Disorder is capable of localizing the wavefunctions, thereby eliminating the option of different parts of the system exchanging energy and equilibrating. It was long believed that disorder induced localisation (better known as Anderson localisation [16]) of single particles would not survive the inclusion of interaction. Recent progress in combining the effects of disorder and interactions have indicated however that it is indeed possible to preserve localization in the presence of interactions: a phenomenon called many-body localization [17]. Effectively, the localization leads to 'emergent integrability' [18, 19]: locally conserved quantities that are conserved not because of symmetries or other conservation laws, but because localization prevents them from changing.

1.3. OUT-OF-EQUILIBRIUM SYSTEMS

All of the above introductions have considered the world of equilibrium physics. Equilibrium physics takes place after the system has settled into a steady state, which typically happens on long timescales. Instead, in this thesis we will mostly be concerned with non-equilibrium physics. In a system that is out-of-equilibrium, it is possible to observe physical phenomena that have no equilibrium counterpart. In order to sustain such physical phenomena, the system is typically driven out-of-equilibrium explicitly. We will not consider such exclicitly driven systems, so that eventually the systems we discuss relax to equilibrium. Rather, we will be concerned with the transient behaviour of the system, i.e. we will study what happens on the way to equilibrium. In particular, we will be discussing two such scenarios.

In Chapter 4, we study the relaxation of initial states that are topologically non-trivial when coupled to external baths. Since we discuss one dimensional systems, these are of the symmetry protected topological type. The coupling to the baths is done via a master equation in Lindblad form [20], meaning that we switch to the language of density matrices [21]. The use of open system theory in this context has two main points. First, we do not wish to describe the near infinity of bath degrees of freedom explicitly. Using the open system formalism to reduce this complexity down to an effective master equation is hence a very practical choice. The second point is that in reality no physical system is perfectly isolated [22]. An external bath is always present (in the very least, the very act of measuring the system requires the 'opening up' to the environment), and so the study of physics in the presence of baths it is rather a necessity than a choice. Of course for observing the physics that would have occured without the bath present, the question is on what timescale the bath has an influence on the system. For the approach to equilibrium however, we will always encounter this timescale.

The second scenario in which we study the transient behaviour is detailed in Chapter 5. Instead of using open systems theory to describe the coupling to a bath, here we consider a small system that we interpret as the bath explicitly. Whereas the master equation in Lindblad form is obtained using various approximations, such as weak-coupling and a

bath with no memory, the approach in Chapter 5 is not based on approximations. As the environment we will consider a disordered and interacting spin-chain, while the system itself consists of a single spin that couples to it. We time-evolve the combined system and environment as a whole, but observe the relaxation of the single spin towards equilibrium. Since the combined system is closed and we consider unitary time evolution, the spin never truly relaxes (due to recurrences). When averaging over different realizations of disorder however, we learn essential features about the environment from only observing the single spin.

1.4. MACHINE LEARNING

Slightly orthogonal to the other topics discussed so far, is the topic of machine learning [23]. Machine learning refers to computer algorithms that make predictions about data, and know how to improve these predictions based on new data they have at their disposal. On the one hand, the 'learning' aspect may be simply regarded as numerical optimization. On the other hand, the machine learns to perform a task that it was *not explicitly programmed to do*. Since nowadays more and more data is availble via e.g. numerical simulations (think of Monte Carlo methods), it makes sense to use machine learning to try and extract features that may otherwise have gone unnoticed by humans.

Machine learning comes in two main variants, namely that of supervised and unsupervised learning. In the latter, the machine is given *only* the data and must try to extract features. In the former, the data has to be prepared in such a way that the inputs are all assigned correct output labels [24]. The machine algorithm then has to optimize its parameters to correctly predict the given output labels, after which one may try to apply it to data it has not seen before.

Often, unsupervised learning is preferred over the supervised method. Mostly this is because the unsupervised methods such as regression or principal component analysis can be more easily understood and interpreted. Whereas performance can arguably be better with supervised methods such as neural networks, in practice one often has no idea *how* or *what* such a network learned. Regardless, we will apply a supervised learning method based on artificial neural networks to try and predict phase transitions from entanglement properties of wavefunctions. Details on this are described in Chapter 3.

Especially the topic of neural networks has seen a recent surge in interest, mainly because it has become possible to train so called 'deep' neural networks. The structure of neural networks is explained in Chapter 3, but suffice it to say that deep neural networks have considerably more parameters to optimize than their shallow cousins. When neural networks where just introduced in the 1950's and 1960's they were expected to revolutionize the world. This promise never happened. In the 1980's more and more data became available, and the popularity of neural networks rose again. It was due to the sheer amount of data available that training them became feasible. Their performance in more 'human' tasks as for example image or speech recognition stayed behind. It was

clear that more complicated networks could do better in such situations due to their ability to form increasingly abstract concepts. The problem was that training these networks simply did not work because convergence was extremely slow.

Today, training of deep neural networks has become possible due to advances in hardware. Most notably, the increase in computational speed for floating point operations in graphical processing units (GPUs)[3] has lead to deep networks making it to the forefront of performance.

[3] This increase in hardware capability is not driven by science, rather it is driven by increasing demands on video games.

II

EQUILIBRIUM

2

AKLT, MPS AND SPT

*We introduce the concepts of symmetry protected topological states, matrix product states and projective symmetry representations, all on the example of the AKLT state. The AKLT state is the groundstate of an exactly solvable spin-*1 *model. It was the first case to support the conjecture by Haldane that integer spin systems should be gapped. The following chapters will make heavy use of the ideas introduced here.*

2.1. MOTIVATION

Spin systems have been studied for about a decennium [25], and are still at the frontier of modern-day physics research. In the next subsection we will provide a short overview of the history of spin chains, and highlight the importance of the so-called AKLT state. Apart from its historical importance, the AKLT state really provides an excellent starting point for introducing multiple concepts.

First, it serves as the simplest and most intuitive case for topologically non-trivial states in one dimension. In fact, the AKLT state is part of a larger class of states belonging to the non-trivial 'Haldane-phase'. These states will turn out to be governed crucially by the presence of symmetries, which is also demonstrated in this chapter.

Second, the AKLT state provides an introduction to matrix product states (MPS). This concept will be built upon heavily in the coming chapters. On the one hand, a matrix product state representation gives direct access to the Schmidt decomposition of the wavefunction. The Schmidt values (which we will introduce later in this Chapter in the form of the entanglement spectrum) that one obtains from such a decomposition will provide the data for the machine learning techniques we will discuss in Chapter 3. On the other hand, the MPS formalism may be extended to open systems so that we can simulate density matrices. This in turn allows us to consider topologically non-trivial states coupled to an environment in Chapter 4.

2.1.1. HISTORY

Especially the history of one dimensional (1D) spin chains is interesting, since the dominant spin-wave theory that was developed in the 1950s [26, 27] didn't apply well to it. For higher dimensions, spin-wave theory made clear predictions about systems having both *long-range order* and *gapless modes*. The exact solution for spin-1/2 chains by Bethe [28] showed that gapless excitations indeed exist in 1D too, but long-range order does not occur in 1D by virtue of the Mermin-Wagner theorem.

The absence of long-range order in 1D lead to the investigation of the possibly that the groundstate was in a Kosterlitz-Thouless (KT) phase. The KT phase is characterized by correlations that decay following a power-law behaviour, and are hence not truly long-ranged but rather 'quasi' long-ranged [29]. The Kosterlitz-Thouless phase indeed also hosts gapless excitations, but a small inconsistency remained. The systems in 1D known to exhibit Kosterlitz-Thouless phases all possessed an SO(2) symmetry, whereas the spin systems under consideration such as the Heisenberg model where invariant under a larger SO(3) symmetry group.

On top of the confusion, in 1983 Haldane conjectured on general grounds that only models with half odd-integer spin ($s = 1/2, 3/2, \ldots$) have a gapless groundstate whereas the integer spin variants have a finite gap [30, 31]. Numerical evidence of this conjecture was soon gathered [32–34], but analytical progress was difficult. This was not in the least

due to the fact that Haldane's arguments were based on topological terms in quantum field theories; a mathematical machinery that was not in wide use yet.

The next big analytical development around Haldane's conjecture was the extension of the Lieb, Schultz and Mattis (LSM) proof [35] from spin-1/2 systems to arbitrary half odd-integer spins. The theorem states that either the groundstate has gapless excitations or the groundstate is degenerate. In addition, it was possible to show that the proof *fails* for integer spin [36]. A proof, or indeed any good examples of gapped integer spins was still lacking.

In 1987 Affleck, Kennedy, Lieb and Tasaki (AKLT) managed to explicitly construct an exactly solvable spin-1 model and its groundstate, now known as the AKLT model and AKLT state respectively [37, 38]. The construction of this state based heavily on the concept of valence-bond solid (VBS) states that were already encountered as the groundstate of a known spin-1/2 model called the Majumdar-Ghosh chain [39]. A valence-bond is nothing but the pairing of two spin-1/2's into a singlet state. In the Majumdar-Ghosh groundstate, two out of every set of three neighboring spins is paired into a valence-bond. That means no set of three spins can ever be in a spin-3/2 state. A Hamiltonian that projects every set of three spins onto the spin-3/2 subspace will hence annihilate this state. Since a projector only has eigenvalues 0 or 1, this state is the one with the lowest possible eigenvalue and hence the groundstate. The realisation of the AKLT paper was that this principle can be extended to higher spins, as we will demonstrate in more detail below.

In the following section, we introduce the AKLT model and its groundstate. Contrary to what one usually encounters in quantum mechanics, in this case we start by constructing an explicit spin-1 state for which we can write down a Hamiltonian to which it is the groundstate, i.e. a parent Hamiltonian. Next we analyze the AKLT state in terms of its non-local order parameter, leading us to a small excursion into symmetries and a discussion on the concept of symmetry protected topological states. From there, it is a small step to generalize all of the encountered concepts and build a full picture of matrix product states (MPS). The concept of MPS can be introduced in various different ways, and one does not need to know about the AKLT state nor about other symmetry protected topological states. Such different viewpoints on MPS can be found in Appendix A, which can be read as a fully independent introduction to MPS.

2.2. AKLT STATE

We start the construction of the AKLT state by considering a set of $2L$ spin-1/2's that we index by a site label $i = 1 \ldots 2L$. We pair up the spins on sites $2i$ and $2i+1$ into of valence-bonds (singlets) as is shown by the lines connecting the circles in in Fig. 2.1. This construction leaves the spins on sites $i = 1$ and $i = 2L$ unpaired for open boundary conditions, but they too can be paired if we consider periodic boundaries. At this point, we have a set of either $L-1$ (open boundaries) or L (periodic boundaries) disconnected valence-bonds. For convenience we stick to the case of periodic boundaries, although

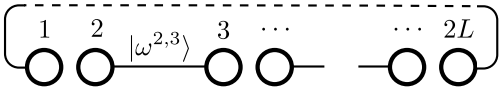

Figure 2.1: A one dimensional chain of spin-1/2 particles denoted by circles, paired up into singlets (valence-bonds) indicated by solid lines. The valence-bond states are given by the $|\omega^{i,i+1}\rangle$ states of Eq. 2.1. The spins at the boundaries are left unpaired for a chain with open boundary conditions, but are paired (dashed line) for the periodic boundary variant.

with minor modifications the following derivation will still hold for open boundaries.

Let us introduce a basis $|\alpha_1\alpha_2\ldots\alpha_{2L}\rangle$ for the spin-1/2's, where $\alpha_i \in \{1,2\}$ enumerates the possible values of $m_z = \{\frac{1}{2},-\frac{1}{2}\} = \{\uparrow,\downarrow\}$ on site i. To construct a valence-bond $|\omega^{i,i+1}\rangle$ between spins at sites i and $i+1$ we use a matrix $\Omega_{\alpha_i,\alpha_{i+1}}$ such that

$$|\omega^{i,i+1}\rangle = \sum_{\alpha_i,\alpha_{i+1}} \Omega_{\alpha_i,\alpha_{i+1}}|\alpha_i\alpha_{i+1}\rangle = \frac{1}{\sqrt{2}}\Big(|\uparrow_i\downarrow_{i+1}\rangle - |\downarrow_i\uparrow_{i+1}\rangle\Big). \tag{2.1}$$

The state with all of the pairs connected into singlets is then given by

$$|\tilde{\psi}\rangle = \sum_{\alpha_1,\ldots\alpha_{2L}} \Omega_{\alpha_2,\alpha_3}\ldots\Omega_{\alpha_{2L},\alpha_1}|\alpha_1\alpha_2\ldots\alpha_{2L}\rangle \tag{2.2}$$

$$= |\omega^{2,3}\rangle\otimes|\omega^{4,5}\rangle\ldots\otimes|\omega^{2L,1}\rangle. \tag{2.3}$$

By itself, this state is not very interesting. It is, after all, a product state of disconnected singlets. What is more, is that we set out to construct a state describing spin-1 degrees of freedom, while what we have at the moment is a collection of spin-1/2's.

To construct a state for spin-1 particles, let us introduce a set of maps $P^{[i]}$ that act on the spin-1/2's at sites $2i+1$ and $2i+2$, i.e. on the spin-1/2's that do not share a valence-bond. We choose these maps such that they project onto the triplet subspace of these two spins, so that we can regard the combination of them as describing a single spin-1. In other words, the $P^{[i]}$ act as maps from the spaces of the two spin-1/2's to that of a spin-1: $P^{[i]}:\mathbb{C}^2\times\mathbb{C}^2\mapsto\mathbb{C}^3$. Without speciyfing the exact form of the $P^{[i]}$'s yet, our state now has the general form

$$|\psi\rangle = P^{[1]}\otimes P^{[2]}\otimes\ldots\otimes P^{[L]}|\tilde{\psi}\rangle. \tag{2.4}$$

This is graphically demonstrated in Figure 2.2. If for the spin-1 on site i we introduce the notation $|s_i\rangle$ with $s=1,2,3$ enumerating the possible values $m_z=\{+1,0,-1\}=\{|+\rangle,|0\rangle,|-\rangle\}$ respectively, we want to make the identification

$$|\uparrow\uparrow\rangle\rightarrow|+\rangle,\quad \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle+|\downarrow\uparrow\rangle)\rightarrow|0\rangle,\quad\text{and}\quad|\downarrow\downarrow\rangle\rightarrow|-\rangle, \tag{2.5}$$

using which a given $P^{[i]}$ can be explicitly written as

$$P^{[i]} = \sum_{\alpha,\beta=1}^{2}\sum_{s=1}^{3} M^{[i],s_i}_{\alpha_i,\alpha_{i+1}}|s_i\rangle\langle\alpha_i\alpha_{i+1}|. \tag{2.6}$$

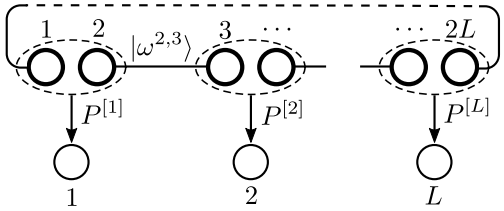

Figure 2.2: The maps $P^{[i]} : \mathbb{C}^2 \times \mathbb{C}^2 \mapsto \mathbb{C}^3$ map two auxiliary spin-1/2 particles to their triplet subspace, thereby identifying it with a $spin-1$ particle.

The $M^{[i],s_i}$ represent a set of 3 (one for every possible value of s_i) matrices of size 2×2 associated to site i. Via the identifications in Eq. 2.5 we can write down the M matrices directly:

$$M^{+} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad M^{0} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad M^{-} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \tag{2.7}$$

We have left out the site indices for convenience, even though for the current case translational invariance already means all of the M matrices are identical.

By substituting the projectors $P^{[i]}$ of Eq. 2.6 into Eq. 2.4, we arrive at the final form of the wavefunction in the spin-1 basis:

$$\begin{aligned}
|\psi\rangle &= P^{[1]} \otimes \ldots \otimes P^{[L]} \left(\sum_{\alpha_1 \ldots \alpha_L} \Omega_{\alpha_2,\alpha_3} \ldots \Omega_{\alpha_{2L},\alpha_1} \right) | \, \alpha_1 \ldots \alpha_{2L} \rangle \\
&= \sum_{s_1 \ldots s_L} \sum_{\substack{\alpha_1 \ldots \alpha_L \\ \beta_1 \ldots \beta_L}} M^{[1],s_1}_{\beta_1 \beta_2} \ldots M^{[L],s_L}_{\beta_{2L-1} \beta_{2L}} \Omega_{\alpha_2,\alpha_3} \ldots \Omega_{\alpha_{2L},\alpha_1} \underbrace{\langle \beta_1 \ldots \beta_{2L} \, | \, \alpha_1 \ldots \alpha_{2L} \rangle}_{= \delta_{\alpha_1,\beta_1} \ldots \delta_{\alpha_{2L},\beta_{2L}}} | s_1 \ldots s_L \rangle \\
&= \sum_{s_1 \ldots s_L} \sum_{\alpha_1 \ldots \alpha_L} M^{[1],s_1}_{\alpha_1 \alpha_2} \Omega_{\alpha_2,\alpha_3} \ldots M^{[L],s_L}_{\beta_{2L-1} \beta_{2L}} \Omega_{\alpha_{2L},\alpha_1} | s_1 \ldots s_L \rangle \\
&= \sum_{s_1, s_2 \ldots, s_L} \mathrm{Tr} \left(M^{[1],s_1} \cdot \Omega \cdot \ldots \cdot M^{[L],s_L} \cdot \Omega \right) | s_1 \ldots s_L \rangle \\
&= \sum_{s_1, s_2 \ldots, s_L} \mathrm{Tr} \left(\tilde{M}^{[1],s_1} \cdot \ldots \cdot \tilde{M}^{[L],s_L} \right) | s_1 \ldots s_L \rangle .
\end{aligned} \tag{2.8}$$

We have defined new matrices $\tilde{M} = M \cdot \Omega$, and have restored the site indices for notational consisitency. As a final tweak, notice that we may introduce a unitary transformation $M \to UMU^{\dagger}$, that leaves the trace unchanged. We can exploit this freedom to impose what is known as 'canonical form' on the M matrices, having them satisfy the relation (with $\mathbb{1}$ being the identity matrix)

$$\sum_{s} \tilde{M}^{s} \left(\tilde{M}^{s} \right)^{\dagger} = \mathbb{1}. \tag{2.9}$$

This particular form is not necessary for the analytical form of the AKLT state. However, having the matrices obey this equation has both drastic consequences for numerical

simulations as well as for their interpretation in relation to entanglement, which we will come to discuss later. The final form of the $\tilde{M}$ matrices obeying Eq. 2.9 is then given by:

$$\tilde{M}^+ = \frac{1}{\sqrt{3}}\begin{pmatrix} 0 & \sqrt{2} \\ 0 & 0 \end{pmatrix}, \quad \tilde{M}^0 = \frac{1}{\sqrt{3}}\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \tilde{M}^- = \frac{1}{\sqrt{3}}\begin{pmatrix} 0 & 0 \\ -\sqrt{2} & 0 \end{pmatrix}. \tag{2.10}$$

In summary, what we have achieved at this point is the representation of a spin-1 state that is built fully from valence-bonds of spin-1/2 auxiliary spins, written in a compact form of Eq. 2.8. This representation is a special case of a matrix product state (MPS), to which we turn in section 2.3, after analyzing the AKLT state in some more detail.

It is actually possible to write down a Hamiltonian to which this state is the exact ground-state. The construction of this parent Hamiltonian is based on the underlying valence-bonds of the state [40]. Consider two neighboring spin-1 objects and the four auxiliary spins that are built up from. The middle two auxiliary spins are paired up into a singlet, and hence have total spin $S = 0$. That leaves the outermost auxiliary spins to pair up into either a singlet too, or a triplet. *Hence the combination of the four auxiliary spins can have at most total spin* $S = 1$*, whereas the two spin*-1 *objects could have paired into total spin* $S = 2$. We see therefore, that in this particular state two neighboring spin-1's are never in a total $S = 2$ state. That means we can write down a Hamiltonian that is a sum over projectors $\mathbb{P}^{S=2}(S_i, S_{i+1})$ that project neighboring spins into the $S = 2$ sector:

$$H = \sum_{i=1}^{L} \mathbb{P}^{S=2}(S_i, S_{i+1}) \tag{2.11}$$

We are guaranteed that the AKLT state is annihilated by each term in this sum, and since the full Hamiltonian is itself a projector with eigenvalues 0 or 1, this is by construction the lowest eigenvalue we can get. Hence, we have constructed explicitly a Hamiltonian to which Eq. 2.8 is the groundstate. For completeness, we note that the Hamiltonian in Eq. 2.11 can also be written in terms of the spin operators explicitly as (up to a constant shift in energy):

$$H_{\text{AKLT}} = \sum_{i=1}^{L} \mathbf{S}_i \cdot \mathbf{S}_{i+1} + \frac{1}{3}(\mathbf{S}_i \cdot \mathbf{S}_{i+1})^2 . \tag{2.12}$$

This is the form usually encountered for the AKLT Hamiltonian. Lastly, we mention that this strategy of using valence-bonds and constructing Hamiltonians as projectors can be applied to higher dimensional spins, and also in higher dimensions. In fact, as mentioned in the history section, this realisation is what inspired the AKLT work in the first place.

2.2.1. EDGE MODES

So far, we have considered the AKLT model with periodic boundary conditions. In terms of the construction of the AKLT state, this means that the first and last auxiliary spin have been paired up into a singlet. If we had not paired them, we would have been

left with two free spin-1/2 degrees of freedom at the ends of our chain. In this case, the groundstate would have been fourfold degenerate; one for each of the four combinations these end spins can be in. In general, for a finitely sized system, these end spins will hybridize and the groundstate will no longer be degenerate. In the thermodynamic limit however, this hybridization becomes exponentially small and the groundstate becomes unique.

The presence of edge modes themselves is not very special. But in the current case, these edge modes stem from a topological property and are hence very robust. We may change the properties of the projection maps $P^{[i]}$ in the construction, but that will not change the fact that the edge spins remain unpaired. Hence these edge modes are in a sense 'built-in' and this is where they get their robustness from. This is not directly obvious from the Hamiltonian in Eq. 2.12, without realizing that this Hamiltonian is equivalent to the projector version in Eq. 2.11. In fact, the AKLT state is just a single point in what is known as the Haldane phase [41–43]. The Haldane phase is a topologically non-trivial phase that is characterized by the presence of the free spin-1/2 degrees of freedom at the edges. Apart from the edge modes, the Haldane phase can be identified via a non-local order parameter called string order [44]. This non-local order parameter relies on the presence of certain symmetries, a point that we will discuss below.

2.2.2. STRING ORDER

Since the Haldane phase is a topological phase and therefore does not spontaneously break any symmetry that may identify it, there is no *local* order parameter. There is a non-local string order parameter, that was first introduced in Ref. [44] as

$$\langle O^\alpha \rangle = \lim_{|i-j|\to\infty} \langle S_i^\alpha \prod_{k=i+1}^{j-1} e^{i\pi S_k^\alpha} S_j^\alpha \rangle, \tag{2.13}$$

in which $\alpha = x, y, z$ selects the spin axis along which we measure string order. It is good to remark here that some other texts have the product over k run up to j instead of $j-1$. The difference is a minus sign, and is not relevant for our current discussion. The equation for the string order observable $\langle O^\alpha \rangle$ in Eq. 2.13 is not very insightful at first glance, but the non-local order that it detects is actually not very involved. A state has non-zero string order if, after stripping out all the spins that are in the $m_z = 0$ configuration, the state has spins alternating between $m_z = +1$ and $m_z = -1$. Let us check that indeed for the AKLT state this is the case.

The following is a useful picture for understanding the string order in the AKLT state. It shows that, due to the valence-bonds of the auxiliary spins, no two spin-1's next to each other can ever both be in the $m_z = +1$ or $m_z = -1$ state [1]. To show this, we consider two neighboring spin-1's and choose the first to be in the $m_z = +1$ state. The underlying two auxiliary spin-1/2's must then be in the $|\uparrow\uparrow\rangle$ configuration. This puts a constraint on

[1] This is nothing new of course, since we already know that two neighboring spins cannot have a total spin $S = 2$.

the next spin-1, since it's corresponding first auxiliary spin is forced to be in the $|\downarrow\rangle$ state by the valence-bond. That leaves the $|\downarrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ configurations, corresponding to the $m_z = 0$ and $m_z = -1$ states. We see hence that it is not possible, by construction, to have two $m_z = +1$ spins next to each other. The same analysis can be done for starting with a $m_z = -1$ state. It is also clear that the $m_z = 0$ state can be indefinitely repeated. So we may conclude that all of the possible configurations we can construct in this way have string order. Since the AKLT state is such a simple state to write down, this can of course also be computed exactly [37, 38] with the result that the string order parameter along any spin axis α takes the expectation value 4/9. Finally let us remark that although this order parameter is non-local, it can be experimentally measured [45].

2.2.3. SYMMETRIES

Unfortunately, the string order parameter as defined in Eq. 2.13 is not the most general string order and only works if the model has at least a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, corresponding to rotations of the spins by 180 degrees about two orthogonal axes[2]. It is possible however, to introduce generalized string orders that work for other symmetry groups [46]. In general though, the non-vanishing of the generalized string order only indicates a symmetric state and therefore does not distinguish topologically trivial from non-trivial states. The string order operator *can* be constructed in such a way that it vanishes in the trivial state, so that non-zero string order is equivalent to being in a topological phase. Overall, string order can not classify the different topological states: it is possible to explicitly construct non-trivial phases that are not picked up by any choice of generalized string order [42]. To classify the various topological phases in 1D, we will need to discuss (projective) symmetries in more detail.

For the discussion on symmetries, it is useful to introduce the anisotropic spin-1 Heisenberg model in a tranvserse field:

$$H = J \sum_{i=1}^{L} \mathbf{S}_i \cdot \mathbf{S}_{i+1} + B S_i^x + D \left(S_i^z\right)^2 . \tag{2.14}$$

We will fix $J = 1$ in the rest of this discussion, and measure both B and D in units of J. This model has been extensively studied [41, 42] and its groundstate phase diagram is well-known (schematically shown in Fig. 2.3). The phase diagram shows four different phases in total. Two of them are of the spontaneous-symmetry-breaking type (the antiferromagnetic ordering states indicated by *AF-z* and *AF-y*), while the two others do not spontaneously break any symmetries (the Haldane phase and the Large-D phase). The latter two phases are therefore not classifiable using the Landau symmetry breaking paradigm, and are topologically distinct. Since for $D \to \infty$ the groundstate becomes a simple product state of spins in the $m_z = 0$ state, this must be the topologically trivial state. We have already seen that the Haldane phase is topologically non-trivial, characterized by the presence of free spin-1/2's at the edges and non-zero string order. The

[2] This symmetry is equivalent to the dihedral symmetry group D_2, describing a global discrete symmetry of rotations by π around the x, y, and z axes. The two groups are equivalent since the product of π rotations around the x and z axes gives the π rotation about the y axis.

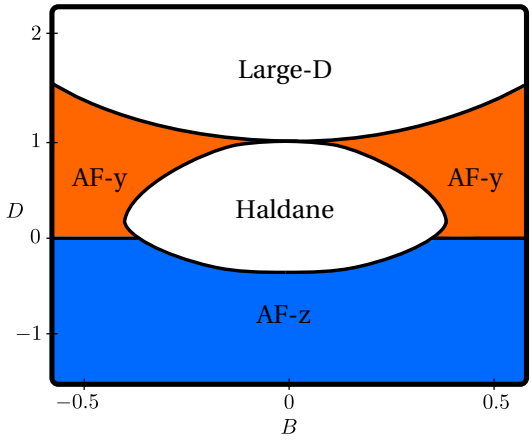

Figure 2.3: Schematic phase diagram of the spin-1 Heisenberg model in a transverse field, as a function of B and D (both in units of J). The model has two symmetry-breaking phases, characterized by non-zero anti-ferromagnetic ordering in the y-direction ('AF-y') and non-zero anti-ferromagnetic ordering in the z-direction ('AF-z'). The central region around the Heisenberg point $(B,D) = (0,0)$ is in the topologically non-trivial Haldane phase, characterized by edge modes and non-zero string order. This phase does not break any symmetries. The Large-D phase does not break any symmetries either, and is hence indistinguishable from the Haldane phase on the level of symmetries. It is connected to the trivial state at $D = \infty$, and does not have edge modes nor string order.

Haldane phase is an example of a symmetry protected topological (SPT) phase: a topologically non-trivial phase that is only non-trivial in the presence of (a set of) certain symmetries. As soon as one of these symmetries gets broken, the topologically non-trivial nature of the state dissapears. Here, that would mean that the distinction between the Large-D phase and the Haldane phase disappears.

The set of symmetries that the Haldane phase and the Large-D phase in Eq. 2.14 are invariant under consists of spatial inversion symmetry, time reversal symmetry and the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry mentioned earlier. What is different in the two phases however, is the representation of these symmetry operations at the *edges* of a chain due to the presence or absence of the edge modes. This holds true also for an artificially introduced edge of an infinite system, i.e. by tracing over one part of the system and looking at the boundary of the left-over semi-infinite half.

The symmetry operators at the edge are in a sense a 'fractionalized' version of the full operators, and correspond to inequivalent projective representations of the symmetry operator [47, 48]. In fact, the projective symmetry representations are powerful enough to provide a full classification of all the possible topologically non-trivial states in 1D. This full classification by projective symmetries works only in 1D, since for 1D the *only* non-trivial phases are SPTs [49]. We review this classification for the pure states in the coming section. A conventient way of doing so is based on matrix product states (MPS), of which we have seen the AKLT state as an example. We therefore turn to a brief discus-

sion on MPS, focussing on the relevant concepts. Both the MPS and projective symmetry representations have a more in depth discussion devoted to them in the form of the Appendices.

2.3. MATRIX PRODUCT STATES

For the purpose of the classification we will consider translationally invariant MPS states, also refered to as iMPS (for *infinite* MPS) [50]. A general such state can be written (c.f. Eq. 2.8)

$$|\psi\rangle = \sum_{s_1,\ldots,s_L} \mathrm{Tr}\left(M^{[1],s_1}\cdots M^{[L],s_L}\right)|s_1\ldots s_L\rangle, \tag{2.15}$$

where the M^{s_i} are matrices of dimension $D\times D$, and where we have introduced basis states $|s\rangle$ with $s=1\ldots d$ enumerating the possible physical states on a given site (d is the dimension of the local Hilbert space). The superscripts on the M are additional labels denoting the site that the M is associated to. For the AKLT state in Eq. 2.8 we found $D=2$, but the groundstate of gapped 1D Hamiltonians can be written in the form of Eq. 2.15 with general *bond dimension* D [51–53]. The interpretation of this wavefunction is that we consider a unit cell of L sites. If we consider a Hamiltonian that is translationally invariant over e.g. just one site, we may still consider the form of Eq. 2.15 and simply set $M^{s_1}=\ldots=M^{s_L}$ without this choice imposing restrictions on further arguments.

2.3.1. DIAGRAMS

One of the advantages of working with matrix product states is that they have a useful diagrammatic representation, eliminating the need of writing and interpreting equations. The $M^{[i]}$ in Eqs. 2.15 and 2.8 are tensors with three indices $\left(M^{[i]}\right)^{s_i}_{\alpha\beta}$, since for every choice of physical index s_i, $M^{[i],s_i}$ is a matrix with two indices. We introduce a diagram for a $M^{[i]}$ as in Fig. 2.4, where each of the 'legs' corresponds to a particular index. For the MPS matrices of a wavefunction $|\psi\rangle$ we use the convention of drawing the physical leg pointing downwards, and upwards for its complex conjugate $\langle\psi|$. The auxiliary indices always face horizontally left and right. Using the diagrams, various operations on

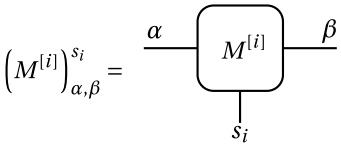

Figure 2.4: Diagram representing an MPS tensor $M^{[i]}$. Such a tensor has three legs, corresponding to the three indices s_i,α and β. For $M^{[i]}$, we use the convention of drawing the physical leg (s_i) downwards, whereas for $\left(M^{[i]}\right)^*$ we would draw it facing upwards.

the MPS matrices can be represented in a straightforward way. An example is shown in

Fig. 2.5 below, in which matrix multiplication is replaced by contraction over the tensor legs. The last important object that we introduce here is the (generalized) *transfer ma-*

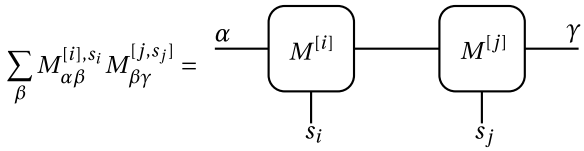

Figure 2.5: Contraction of MPS diagrams corresponds to summing over the respective indices of the involved tensors. In this example, we represent the matrix multiplication of $M^{[i]}$ with $M^{[j]}$ by contracting the corresponding auxiliary legs. Applying a physical operator would be represented by contracting it on the physical leg instead.

trix $\mathbb{T}^{[i]}_X$ shown in Fig. 2.6, where X is some general physical operator. In an equation, the generalized transfer matrix for site i would read

$$\mathbb{T}^{[i]}_X = \sum_{s_i, s_i'} M^{[i],s_i} X_{s_i,s_i'} \left(M^{[i],s_i'}\right)^* . \tag{2.16}$$

Usually when we mention the transfer matrix, we refer to the transfer matrix of the entire MPS unit cell. In other words, we would construct $\mathbb{T} = \mathbb{T}^{[1]}_{X_1} \mathbb{T}^{[2]}_{X_2} \dots \mathbb{T}^{[L]}_{X_L}$. For simplicity though, we will mostly consider the example of an MPS with a unit cell consisting only of a single site. In that case, the transfer matrix for this single site is identical to the full transfer matrix. The transfer matrix will play a special role below in section 2.4 for the

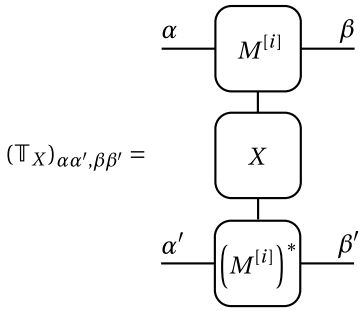

Figure 2.6: The generalized transfer matrix $\mathbb{T}^{[i]}_X$ for site i and general physical operator X.

determination of projective representations of symmetry operators. We will leave the details to Appendix A, but revist shortly the concept of canonical form here. The canoncial form mentioned in Eq. 2.9 can namely be stated as a property of the transfer matrix $\mathbb{T}$ (with the operator X being the identity matrix, and left out for notational convenience): *An MPS is in left-canonical form if the transfer matrix has the identity matrix as its dominant left eigenmatrix with dominant eigenvalue η s.t. $|\eta| = 1$. All other eigenvalues have*

magnitude strictly less than unity[3]. A similar statement holds for the right-canonical form, and in fact it is possible to introduce an elegant notation [54] that automatically guarantees both left- and right-canonical form. The canonical form leads to the diagram shown below in Fig. 2.7. The canonical form drastically reduces the (numerical) com-

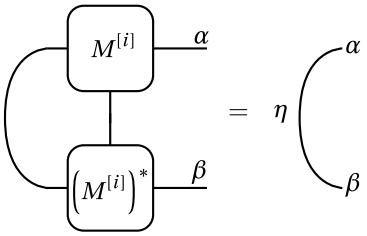

Figure 2.7: The (left-)canonical form of an MPS means the transfer matrix (represented here using the matrices $M^{[i]}$ only), representing that it has the identity matrix as a left eigenmatrix with eigenvalue η. If the MPS is in canonical form, $|\eta| = 1$ and is the only eigenvalue with unit modulus.

plexity for the calculation of expectation values. From now on, we will always assume canonical form for the MPS.

2.4. PROJECTIVE SYMMETRY REPRESENTATIONS

Let us consider a system that is invariant w.r.t. a symmetry group $\mathcal{G}$, and has a ground-state $|\psi\rangle$ that does not spontaneously break any of the symmetries. We will restrict ourselves to the discussion of *local* symmetries, meaning that for a local symmetry operator $R_g \in \mathcal{G}$, the groundstate wavefunction $|\psi\rangle$ should obey

$$|\tilde{\psi}\rangle = R_g \otimes \cdots \otimes R_g |\psi\rangle = e^{iL\phi_g}|\psi\rangle, \tag{2.17}$$

where the R_g each act on a site of the system, and with ϕ_g a phase factor that may depend on the specific group element. We now ask what consequences this has for the MPS matrices, given that we write the groundstate wavefunction as an MPS.

The symmetry operators R_g have a representation in matrix form when they act on the basis states $|s_i\rangle$:

$$|s_i'\rangle = R_g|s_i\rangle = \sum_{s_i} \left(R_g\right)_{s_i,s_i'} |s_i\rangle. \tag{2.18}$$

This expression may be directly substituted in Eq. 2.17 with $|\psi\rangle$ written in an MPS form

[3] If the state is a cat-state, i.e. a superposition of (two) macroscopically polarized states, there is more than on eigenvalue with unit magnitude [46].

as in Eq. 2.15 to obtain:

$$\begin{aligned} R_g \otimes \cdots R_g |\psi\rangle &= \sum_{s_1,\ldots,s_L} \mathrm{Tr}\left(M^{s_1}\ldots M^{s_L}\right) R_g \otimes \cdots R_g |s_1 \ldots s_L\rangle, \\ &= \sum_{\substack{s_1,\ldots,s_L \\ s'_1,\ldots,s'_L}} \mathrm{Tr}\left(M^{s_1}\ldots M^{s_L}\right) \left(R_g\right)_{s_1 s'_1} \left(R_g\right)_{s_L s'_L} |s'_1 \ldots s'_L\rangle, \\ &= \sum_{\substack{s_1,\ldots,s_L \\ s'_1,\ldots,s'_L}} \mathrm{Tr}\left(M^{s_1}\left(R_g\right)_{s_1 s'_1} \ldots M^{s_L}\left(R_g\right)_{s_L s'_L}\right) |s'_1 \ldots s'_L\rangle, \\ &= \sum_{s'_1,\ldots,s'_L} \mathrm{Tr}\left(\tilde{M}^{s'_1}\ldots \tilde{M}^{s'_L}\right) |s'_1 \ldots s'_L\rangle. \end{aligned} \tag{2.19}$$

In the last line we introduced the transformation for the MPS matrices:

$$\tilde{M}^{s'_i} = \sum_{s_i} M^{s_i} \left(R_g\right)_{s_i s'_i} \tag{2.20}$$

We know that the new wavefunction is supposed to be invariant under this local symmetry operation, and should be related to the untransformed $|\psi\rangle$ by the phase factor $e^{iL\phi_g}$ of Eq. 2.17. This is fulfilled by introducing the following transformation rule for the MPS matrices:

$$M^{s_i} \rightarrow e^{i\phi_g} U_g M^{s_i} U_g^\dagger. \tag{2.21}$$

Here, U_g is a unitary matrix representing the symmetry operation R_g directly on the level of the MPS matrices. As we now vary the element R_g over the whole group we obtain a set of phases ϕ_g and matrices U_g, each forming a representation of the group. In particular, the matrices U_g form a D-dimensional projective representation, since they are $D \times D$ matrices and we may multiply them by add an arbitrary phase factor leaving Eq. 2.21 unchanged.

More importantly for projective symmetry representation we have that if the group elements obey $R_g R_{g'} = R_{gg'}$, the matrices obey $U_g U_{g'} = e^{i\phi(g,g')} U_{gg'}$. In fact, knowledge of the phase factors $\phi(g,g')$ for the elements of the symmetry group provides a full classification and identification of one dimensional topological phases [48, 49, 55].

As an example, let us consider the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry group for the AKLT Hamiltonian in Eq. 2.12. The four elements of this group are $\{\mathbb{1}, R_x, R_z, R_{xz}\}$. We denote the projective representations of this symmetry group via U_x, U_z and U_{xz}, and in particular we wonder about the phase factor in $U_x U_z = e^{i\phi(x,z)} U_{xz}$. As we will see below (or in Appendix B), it turns out that $\phi(x,z) = \pm 1$. This means the model has two different projective representations, one of which we will call topologically trivial and the other topologically non-trivial. Notice that although the information on the topological states is in general contained in the phase factors $\phi(x,z)$, in this case that directly translates to the U_x and U_z matrices either commuting or anti-commuting.

This raises the question of how to obtain the U_g matrices corresponding to symmetry elements R_g. It is here that the transfer matrix of Eq. 2.16 plays an important role. We can

easily prove diagramatically that the matrix U_g is directly obtained as the eigenmatrix of $\mathbb{T}_{R_g}$, as is described in Fig. 2.8. In the following section, we will explicitly construct the U_g matrices and examine the phases for the AKLT model.

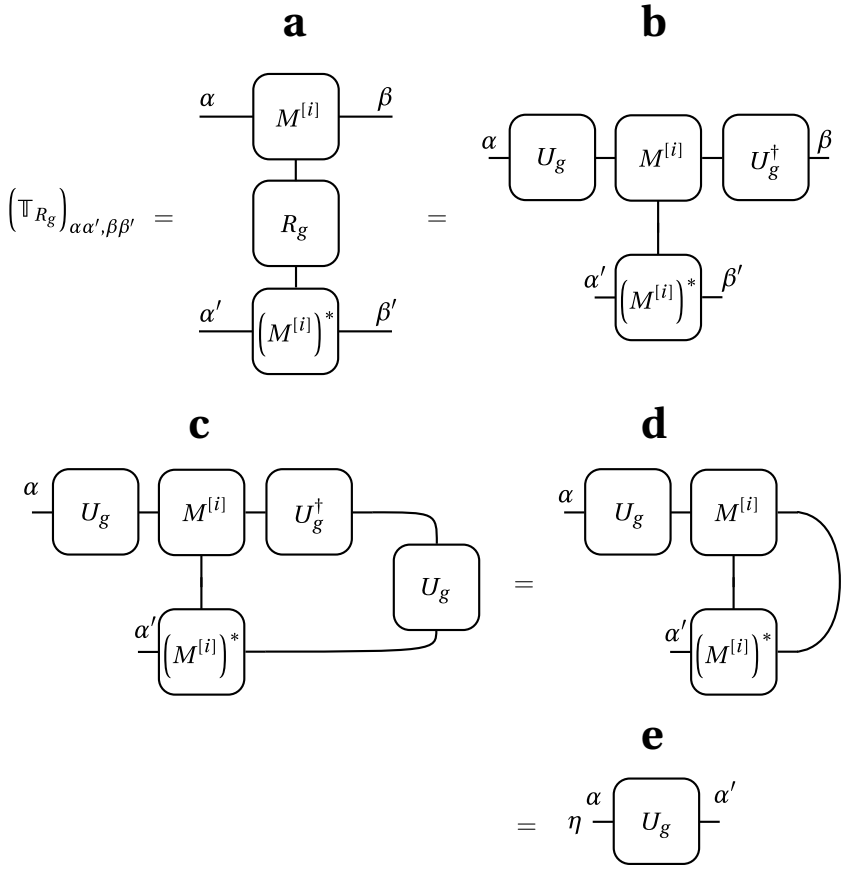

Figure 2.8: The transfer matrix $\mathbb{T}_{R_g}$ in diagrammatic form is shown in part **a**, and simplifies to the diagram shown in **b** after using Eq. 2.21. We have left out the phase factor for convenience. In parts **c**, **d** and **e** we show that the transfer matrix has U_g as its right eigenmatrix. After contracting the transfer matrix with U_g in **c**, we recognize that $U_g^\dagger U_g$ is proportional to the identity matrix (U_g is a unitary matrix). In part **d** we then make use of the canonical form of the M matrices shown in Fig. 2.7, finally arriving at the desired result that U_g is a right eigenmatrix of $\mathbb{T}_{R_g}$. If the eigenvalue η does not have unit modulus, the wavefunction represented by the MPS was not symmetric under the application of R_g and the resulting U_g should be disregarded.

2.4.1. PROJECTIVE SYMMETRIES ON THE AKLT STATE

Since the AKLT matrices are only 2×2 matrices, we have no need for the full machinery of the transfer matrix. Instead, knowledge of the properties of the Pauli matrices σ_x, σ_y

and σ_z is sufficient. The MPS matrices for the AKLT states in Eq. 2.10 can be represented using the Pauli matrices as

$$\tilde{M}^+ = \sqrt{\frac{1}{6}}(i\sigma_y + \sigma_x), \quad \tilde{M}^0 = -\frac{1}{\sqrt{3}}\sigma_z, \quad \tilde{M}^- = \frac{1}{\sqrt{6}}(i\sigma_y - \sigma_x). \tag{2.22}$$

Next, we need the symmetry operators R_x and R_z in matrix form so that we can compute the transformed M matrices as in Eq. 2.20. Rotation operators can be computed from the knowledge of the infinitesimal generators, which for the spin-1 case are the generators S_x, S_y and S_z of $SU(3)$:

$$S_x = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad S_y = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & -i & 0 \\ i & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \quad S_z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{2.23}$$

Rotations by π around the x and z axes are then straightforwardly computed as

$$R_x = e^{i\pi S_x} = \begin{pmatrix} 0 & 0 & -1 \\ 0 & -1 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad R_z = e^{i\pi S_z} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \tag{2.24}$$

After applying these operators to the M matrices (c.f. Eq. 2.20) one may recognize the transformations

$$\tilde{M}^{s_i'} = \sum_{s_i} M^{s_i} (R_x)_{s_i s_i'} = \sigma_x M^{s_i'} \sigma_x \tag{2.25}$$

$$\tilde{M}^{s_i'} = \sum_{s_i} M^{s_i} (R_z)_{s_i s_i'} = \sigma_z M^{s_i'} \sigma_z \tag{2.26}$$

$$\tilde{M}^{s_i'} = \sum_{s_i} M^{s_i} (R_x R_z)_{s_i s_i'} = \sigma_y M^{s_i'} \sigma_y, \tag{2.27}$$

from which we can readily identify the projective representation matrices:

$$\{U_x, U_z, U_{xz}\} = \{\sigma_x, \sigma_z, i\sigma_y\}.$$

The extra factor i in the representation U_{xz} is a choice out of convenience, allowed by the phase degree of freedom alluded to for Eq. 2.21. Most importantly, we may now conclude that

$$U_x U_z + U_z U_x = 0,$$

meaning U_x and U_z anti-commute. We may show on general grounds (in Appendix B) that this corresponds to the topologically non-trivial phase of the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry group[4]. This again confirms that the AKLT state is topologically non-trivial. For any other state in the Haldane phase, we may obtain the U_x and U_z matrices from the transfer matrix approach in Fig. 2.8 and ask whether they commute or anti-commute.

Finally, in the next section we consider one more method of identifying and characterizing (topological) phases of matter.

[4] To be more precise, the set of phases define a 2-cocycle that corresponds to the non-trivial element of the second cohomology group $H^2(\mathbb{Z}_2 \times \mathbb{Z}_2, U(1)) = \mathbb{Z}_2$. This is explained in Appendix B

2.5. ENTANGLEMENT SPECTRUM

Over the last few years, entanglement has started to play a more and more prominent role in condensed matter physics [56]. For our purposes, we focus on the main concepts of entanglement spectrum and entanglement entropy. Both of these are obtained from (and defined for) a bipartitioning of the system. The entanglement spectrum was first used suggested as a powerful tool in Ref. [57], and contains more information than the entanglement entropy. The latter can be obtained from the spectrum as a simple sum.

Let us perform an arbitrary bipartitioning of the system into parts A and B. Usually we consider a bipartitioning in which we spatially cut the system in half, so that A and B may be identified as the left and right half respectively. It is important to keep in mind that this is by no means the only (or best?) option. We could have chosen the partitioning to be between extensive sections of the bulk [58], have made the partitioning in momentum space [59, 60] or indeed even have chosen random partitionings [61].

In either case, to obtain the entanglement spectrum we first compute the full density matrix of the joint system $A \cup B$, denoted $\rho_{A\cup B}$. We may then trace out either part A or B to obtain the reduced density matrix for the part that is left over. We will choose trace out part B, and obtain

$$\rho_A = \mathrm{Tr}_B \rho_{A\cup B}.$$

The entanglement spectrum[5] is defined as the collection of values $-\ln\lambda_i$, where the λ_i are the eigenvalues of ρ_A. Since the λ_i are eigenvalues of a physical density matrix, they are constrained by $\sum_i \lambda_i = 1$ and are all real and positive. Had we traced out part A to obtain ρ_B, we would have found an identical set of eigenvalues, and so the choice is arbitrary. In addition, the entanglement entropy is defined as $S = -\mathrm{Tr}\rho_A \ln\rho_A$ and may be computed from the entanglement spectrum as $S = -\sum_i \lambda_i \ln\lambda_i$.

The eigenvalues λ_i may also be directly computed from a Schmidt decomposition of the wavefunction, as described in Appendix A. The Appendix also shows that the MPS representation of a wavefunction allows for direct access to the entanglement spectrum from the M matrices, without having to perform a partial trace.

The entanglement spectrum may be used to distinguish the Haldane phase from the Large-D phase. Like the projective symmetry representations it relies only on the symmetry group of the model, and is therefore more robust than the string order parameter. Topological phases are characterized by the structure of *degeneracies* of the levels [48]. This can be deduced from the way the symmetry operations are represented on the reduced density matrices.

This is most elegantly demonstrated by considering spatial inversion symmetry. The model in Eq. 2.14 possesses this symmetry, and will serve as a good example. We decom-

[5] That this is refered to as a spectrum comes from interpreting the reduced density matrix as the logarithm of the entanglement Hamiltonian: $\rho_A = \exp\left(-H_{\mathrm{ent}}\right)$. This can always be done by virtue of ρ_A being a Hermitian matrix.

pose the $M^{[i],s_i}$ matrices of the matrix product state into a product $M^{[i],s_i} = \Gamma^{[i],s_i}\Lambda^{[i]}$, where $\Lambda^{[i]}$ is diagonal and contains the singular values of $M^{[i]}$. This notation was introduced first by Vidal, and allows for expressing the left- and right-canonical form in a more compact way. Namely, the MPS is in canonical form if (compare Eq. 2.9)

$$\sum_{s_i} \Gamma^{s_i}\left(\Lambda^{[i]}\right)^2\left(\Gamma^{s_i}\right)^\dagger = \sum_{s_i}\left(\Gamma^{s_i}\right)^\dagger\left(\Lambda^{[i]}\right)^2\Gamma^{s_i} = \mathbb{1}. \tag{2.28}$$

If these conditions are fulfilled, then for any local unitary operator U one may prove [46, 62] that $[U,\Lambda] = 0$[6].

Let us get back to the inversion symmetry. A spatial inversion on the level of the MPS matrices means that $\Gamma \to \Gamma^T$, and we may find a representation of this in terms of a unitary matrix U_I such that

$$\Gamma^T = e^{i\theta}U_I\Gamma U_I^\dagger. \tag{2.30}$$

If we repeat this equation, by substituting it back into itself, we find that the matrix $U_I U_I^*$ is an eigenmatrix of the transfermatrix with eigenvalue $e^{-2i\theta}$. By our assumption of canonical form, that means $\theta = 0$ or π (or other integer multiples of pi) and $U_I U_I^* = e^{-i\phi}\mathbb{1}$. From this last equation we find that $U_I^T = e^{i\phi}U_I$, and with ϕ also either being 0 or π (from repeating the equation). This means that U_I is either a symmetric or anti-symmetric matrix. This difference has an important consequence. Namely, if U_I is antisymmetric the singular values in Λ must come in degenerate pairs, since it commutes with U_I.

More importantly, the multiplicity must be even. This is straightforwardly shown by considering the subblock of U_I corresponding to such a degenerate pair of eigenvalues. Let us say call the multiplicity of this pair α, and hence the sublock U_I^α is an $\alpha\times\alpha$ matrix. We consider the determinant $\det U_I^\alpha = \det\left(U_I^\alpha\right)^T = \det\left(-U_I^\alpha\right) = (-1)^\alpha \det U_I^\alpha$. Since U_I is a unitary matrix it has a non-zero determinant and hence we must have α even s.t. $(-1)^\alpha = 1$.

Even though we have done this analysis using the inversion symmetry of the Haldane phase, we could equally well have obtained the same result using the $\mathbb{Z}_2\times\mathbb{Z}_2$ symmetry we have discussed before, with only minor considerations [47].

[6] A sketch of the proof goes as follows. We may interpret the transfer matrix $\mathbb{T}_X$ of Eq. 2.16 as a linear map. If the state is in canonical form, this map is completely positive and trace preserving. The spectral radius of this map is less than or equal to 1, with equality holding only if a unitary matrix U exists for which Eq. 2.21 holds. If the spectral radius were larger than unity, the wavefunction would not be normalizable in the thermodynamic limit. The singular values in the diagonal matrix Λ are the eigenvalues of $MM^\dagger$, meaning

$$\begin{aligned} MM^\dagger &= \left(e^{i\theta}UMU^\dagger\right)\left(e^{i\theta}UMU^\dagger\right)^\dagger \\ &= UMU^\dagger UM^\dagger U^\dagger \\ &= UMM^\dagger U^\dagger. \end{aligned} \tag{2.29}$$

In other words, $MM^\dagger$ commutes with U implying that $[\Lambda, U] = 0$.

So the entanglement spectrum for the Haldane phase comes in degenerate pairs of even mutliplicity, whereas that of the large-D phase will have non-degenerate and odd-multiplicity degenerate pairs. This is clearly demonstrated in Figure 2.9, showing part of the entanglement spectra obtained from MPS methods of computing the groundstate. The en-

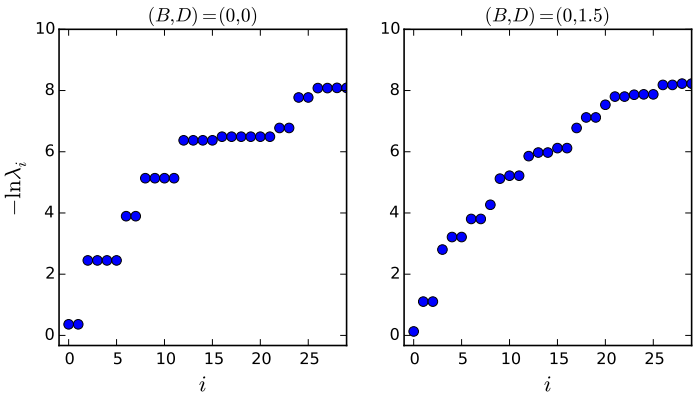

Figure 2.9: The entanglement spectrum for the same two points in the phase diagram of Fig. 2.3. In the Haldane phase, the entanglement spectrum shows a clear pattern of degenerate pairs of even multiplicity. The Large-D phase however, additionally shows non-degenerate pairs and pairs that come with odd-multiplicity degeneracy. This is a consequence of the way the reduced density matrix in each of the phases transforms under the set of symmetries of the Hamiltonian.

tanglement spectrum allows us to clearly distinguish the Haldane phase and the large-D phase. Generally, different SPT phases have different patterns of degeneracies [48].

The entanglement spectrum, although basis dependent, has been conjectured to also contain information on the physical state of the system (see Ref [63] and the references therein). This makes it a very good candidate for 'compressing' the information in an exponentially large wavefunction down to a smaller set of numbers and their structure. In the next chapter, we will study machine learning of phase transitions, and the entanglement spectrum provides an excellent representation of the wavefunctions to learn from.

3

MACHINE LEARNING

In this chapter we review some of the recent progress in applying machine learning to condensed matter physics. We then develop our own approach, using a neural network to identify phase transitions from the entanglement spectra of three different models. The method we develop here is more general than either the entanglement spectrum or condensed matter physics, and may find use in other situations as well.

3.1. GENERAL

Typically, a programmer writes a program for a machine to follow explicitly. For each of the various scenario's, the programmer has to make sure an apropriate output is generated by the program. If there is a mistake in the output, that is solely due to an error on the part of the programmer. The machine itself will not *learn* from its mistake. This is exactly the important difference between regular algorithms and machine learning algorithms. In the latter, the machine will learn to produce a correct output without ever having been explicitly programmed to do so.

The field of machine learning has evolved mainly from the larger encompassing field of artificial intelligence. Tasks such as image recognition, speech recognition and natural language processing have all contributed to what is now refered to as machine learning. The field has a large overlap with probability theory and optimization theory, with both of these topics forming the core of machine learning theory.

Let us make clear up front what we mean when we say a machine has 'learned'. We do this by asking the same question for when a human can be said to have learned something. As an example, we consider a student in a physics course. At the end of the course, we need to design a suitable test, or exam, for assessing whether or not the student has learned about the taught subject. If on the test we were to ask questions about a completely different topic, say Linguistics, the student will most likely not perform very well. Of course, this was not a representative test, since we took the questions from a completely different set of topics. On the other hand, if we were ask the student questions with solutions that were covered explicitly during the lectures, the student will perform very well. This is also not a representative test, since a simple memorization of these question would have lead to the same performance. Instead, we must ask questions that are similar but require the student to apply his/her knowledge to new situations. If he/she performs well, we may conclude that he/she learned about the topic.

This *generalization* skill is arguably the most important demonstration of having learned something, and we can directly carry this notion over to machine learning. Let us imagine that by some procedure we feed the machine some data to be learned, and afterwards ask it to generalize make statements about new data. If the new data is taken from a different underlying distribution than the initial data, we are being unfair. Similarly, if the machine has 'overfit' (i.e. memorized) the input data (i.e. reaching perfect performance), its statements about the new data can not be trusted. Designing suitable learn- and test-data is therefore an important topic that should not be disregarded.

Let us get back to the procedure of feeding the machine data to be learned. We distinguish three main ways of learning. Perhaps the most intuitive, is that of reinforcement learning. In reinforcement learning the learner keeps track of the state of the system, and has an internal table that tells it which action is best taken given the current state. This table is continuously updated by means of a reward function: if the learner chooses an action that leads to a negative reward, it had better not do that again the next time. On the contrary, if its chosen action leads to a desirable outcome it should enforce itself

to try it again.

The other two ways of learning are those of *supervised* and *unsupervised* learning (in fact, reinforcement learning may be seen as a combination of them, as will become clear later). As an example for both, we imagine the task of learning to recognize pictures of objects. In other words, we have a large amount of data samples that we wish to *classify*. Let us say we have N pictures of $M < N$ different objects. In an unsupervised learning approach, we give the learner *only* the data and ask it to try and find a measure by which we can distinguish the inputs. Typical examples that one may think of here are clustering algorithms such as K-Means, or feature extraction algorithms such as principal component analysis (PCA). To demonstrate learning, we make sure that we keep a few of the images hidden. After the machine has e.g. extracted a set of features using PCA, we then ask it to classify these hidden images. By counting how many of these are correctly classified, we have a simple measure of how well the machine has learned.

In supervised learning on the other hand, we *label* the data. Each of our input images is now labeled with what the image actually is (which class it belongs to), and we provide the machine with pairs of input and corresponding correct output. Also in this approach we keep a set of hidden input data separate. After the training stage, we ask the machine to predict the labels for these new samples and compare them with the correct labels. This again gives us a simple measure of the learning of the machine.

It is essential in all of the above methods that a large amount of data is available and that a pattern exists in this data (i.e. that it can be classified in the first place). Large amounts of data have been available in condensed matter physics for a while, but people have only recently begun to try and analyze it using machine learning. This may well be fueled by the fact that machine learning itself has recently received a boost in attention due to breakthroughs in a subfield called 'deep learning'. Machine learning is not new, but the advances in computing power –notably the tremendous increase in GPU hardware power thanks to demands on video games– allow for training models that would have taken infeasibly long 10 years ago.

In this chapter we highlight some of the studies that attempt to use machine learning approaches to condensed matter physics, of which there are examples of both supervised and unsupervised learning. Next we explain a particular machine learning method for supervised learning, namely that of artificial neural networks. Last, we describe our own attempt at using a neural network to classify and extract information of phase transitions from the entanglement spectrum.

3.2. REVIEW

Machine learning has had its reflections on modern condensed matter physics, but the reverse is also true. Concepts such as phases and phase transition have been used in the study of neural networks, and so have tensor networks. In particular, tensor networks provide an efficient way of compression information. In the case of condensed

matter, this usually means an effictive representation of the wavefunction using tensors. This idea can be applied to artificial neural networks too, by effectively compressing the network itself [64] or by introducing a network representation for the wavefunction and variationally optimizing its parameters [65].

The commonality between the fields of machine learning and condensed matter physics is exponential growth. In machine learning, trying to learn one more feature requires exponentially more data[1]. For condensed matter, adding one more site to the system leads to an exponentially larger Hilbert space.

The exponential growth of the Hilbert space means that numerical approaches in condensed matter physics resort to two tricks: *sampling* of the exponentially large space, or effictively *compressing* the data. The matrix product states we have encountered in the previous chapter do exactly the latter, while for example (quantum) Monte Carlo methods use the sampling method. The machine learning community has essentially found the same underlying tricks, but in many different forms. The compression of the data is done in the form of feature extraction, whereas the sampling of the input is typical in artificial neural network approaches.

Let us go back to highlighting how machine learning has already been applied to condensed matter systems. Classical spin systems have provided a good basic training ground for investigating the use of supervised and unsupervised learning methods. In both cases, the input is simply a collection of spin-configurations at different temperatures that were obtained by Monte Carlo methods. In the unsupervised learning approach [66], the configurations were analyzed using principal component analysis (PCA). We will describe PCA in a bit more detail later in this chapter, but it suffices to say at this point that it is a linear method for extracting features. It was shown in Ref. [66] that this approach is capable of extracting order parameters and structure factors directly from the configurations, and is therefore also capable of predicting the transition temperature. Additionally, it allowed the author to find an order parameter for a model in which this was not yet known.

In the supervised learning approach [67], a neural network was trained on the same inputs but this time with an associated label indicating whether the configurations corresponds to above or below the critical temperature T_c. The neural network also learns to distinguish configurations below and above T_c, and toy models were used to argue that the network indeed extracts concepts such as an order parameter to do so.

As a last example, we mention a more fundamental use of neural networks in physics. Namely, we may choose to represent a wavefunction as a neural network in the same sense as we may choose to represent a wavefunction as an MPS. This is the approach taken in Ref. [65], where a particular type of network (a restricted Boltzmann machine) was used as an Ansatz for a variational method of finding the groundstate of a Hamiltonian. This approach was shown to be able to outperform state-of-the-art results using tensor network algorithms for a 2D spin-1/2 Heisenberg model.

[1] This fact is known in the machine learning community as the 'curse of dimensionality'

Motivated by all of the above results, we will study the use of neural networks in identifying phases and phase transitions from the entanglement spectrum. The choice of the entanglement spectrum as an input has many advantages. First and foremost, it is an efficient compression of the otherwise exponentially large wavefunction. Second, there is a large amount of literature available on extracting physical information from it. This means there is good reason to believe that a machine learning approach is capable of extracting these patterns.

In the next section, we first explain what a neural network is. The following sections then deal with its use for identifying phase transitions, part of which is contained in Ref. [68].

3.3. NEURAL NETWORKS

Artificial neural networks were initial modeled after simplifications of real neurons. Physically, such a neuron is a single cell with many dendrites (inputs) and a single axon (output) that connects to the dentrites of other neurons via synapses. A neuron 'fires', i.e. sends an electrical signal down its axon, whenever it receives a large enough stimulus on its dendrites. In essence, it can be thought of as a simple threshold-function: whenever the sum of its inputs is larger than some threshold the neuron provides an output, otherwise it does nothing.

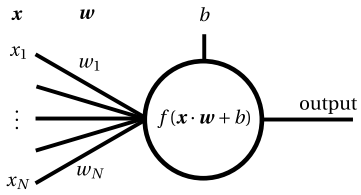

Figure 3.1: The abstraction of a physical neuron into a box with many inputs and a single output. The output is computed from the inputs $\boldsymbol{x}$ by first weighting them with $\boldsymbol{w}$ and adding a bias term b, and then passing the result through an activation function: output $= f(\boldsymbol{x}\cdot\boldsymbol{w}+b)$. This activation function could be a thresholding in the form of a step function, or a more smooth version described by a sigmoid function (see the main text). Notice that the bias parameter b is used to set the threshold.

We model a single such simple neuron as shown in Fig. 3.1. Let us denote the inputs to the neuron by a vector $\boldsymbol{x}$, such that input channel i receives the binary value $x_i \in \{0,1\}$. The inputs may be of different importance (or rather, the physical intensity of input to a dendrite may differ), and so each input channel is associated with a weight vector $\boldsymbol{w}$ representing this. Our simple thresholding neuron can then be mathematically described via

$$\text{Output} = \begin{cases} 0 & \text{if } \boldsymbol{x}\cdot\boldsymbol{w}+b<0 \\ 1 & \text{if } \boldsymbol{x}\cdot\boldsymbol{w}+b\leq 0 \end{cases}, \tag{3.1}$$

where we have introduced the bias variable b determining the threshold. To be precise, what we have now constructed is also known as a *perceptron* [69]. It takes binary values

as its input, and produces a binary value on its output. Such a simple perceptron is already quite a general object capable of computation. For example, taking a perceptron with two inputs with weights $w_1 = w_2 = -1/3$ and a bias $b = 1/2$ is equivalent to a NAND gate. Since the NAND gate is a universal gate (i.e. any other gate can be built from NAND gates only), a *network* of perceptron units is capable of computing just as a normal computer is. In fact, one may prove that a neural network is in principle capable of fitting any function [70].

For practical purposes, we introduce a very similar but slightly different type of neuron. The motivation for this is that we would like the output of the neuron to depend smoothly on its input, instead of jumping discontinuously at the threshold. Hence we will allow the inputs and outputs to be continuous numbers instead of just binary. The neuron we will be using is also known as a sigmoid-neuron, since its activation function is given by a sigmoid:

$$\text{Output} = \sigma(\boldsymbol{x}\cdot\boldsymbol{w} + b) = \frac{1}{1 - e^{\boldsymbol{x}\cdot\boldsymbol{w}+b}}. \tag{3.2}$$

3.3.1. BUILDING THE NETWORK

Let us now construct a neural network with sigmoid neurons as follows. We formally distinguish three types of layers, even though they are all composed of similar neurons. First is the input layer, consisting of neurons whose inputs we disregard and whose outputs we fix to the values of the input $\boldsymbol{x}$. These outputs are fully connected to a layer of neurons called a hidden layer via a matrix of weights. That is, if we have N input neurons and M neurons in the hidden layer, their connections can be described by an $N \times M$ matrix of weights that we denote W_{ij}. Additionally, the hidden layer neurons are associated a vector of biases $\boldsymbol{b}$, so that b_j is the bias for neuron j in the hidden layer. Next, we may add many more hidden layers with varying numbers of neurons and associated weight matrices, until we come to the final output layer. Typically, the output layer consists of only a few neurons, the outputs of which is what constitutes the prediction of the network. How exactly the prediction is encoded in their outputs depends on the representation of the output, which we will come to shortly.

The shape of the network is refered to as the network architecture. Hence a network with an input layer of 50 neurons, two hidden layers of each 100 neurons and an output layer of 2 neurons has an architecture $[50, 100, 100, 2]$. An explicit demonstration of a small network of architecture $[3, 7, 1]$ is shown in Fig. 3.2.

TRAINING THE NETWORK

The goal of training a neural network is now to find a procedure by which we can choose the weights and the biases of the entire network, so that for a given input we manage to predict the correct output. This requires us to discuss the form of the input and output of the network explicitly.

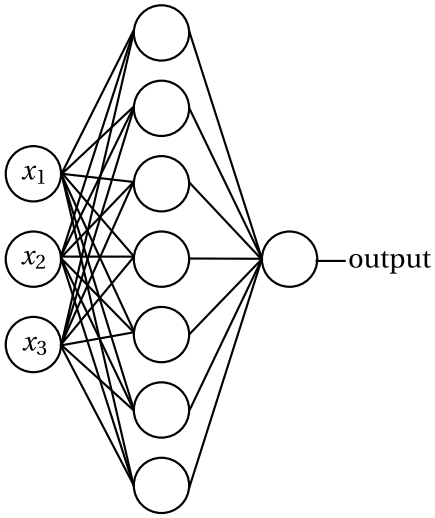

Figure 3.2: A fully connected neural network of architecture [3, 7, 1]. In the input layer, we initialize the neurons to output the values $\boldsymbol{x}$. Notice that we have left out the bias variables for the hidden neurons for simplicity, altough in the actual implementation each of them has its own bias. The connections between the input layer and the hidden layer are described by a weight matrix $\boldsymbol{w}^1$, whereas those of the hidden layer to the output layer by the weight matrix $\boldsymbol{w}^2$ (which in this case due to the single output neuron is actually a vector).

Our data comes in pairs (x, t) with x representing the input data and t representing the correct (i.e. targetted) label. To be concrete, let us think of the input data being 10×10 grayscale pictures of m different objects, and the target t being a number between 1 and m indicating which object it is. The input to the neural network can then be given as a vector $\boldsymbol{x}$ of length 100, representing the intensity of each of the pixels in the image [2]. As an output, we choose to use the popular 'one-hot' encoding. This means that our output layer has m neurons, and we wish to train the network such that neuron j has a large output if the correct target for the input was $t = j$, while the other neurons have small outputs.

Initially, the network is initialized with random weights. We split all of the data we have into two sets that we denote $\mathscr{L}$ (for learning) and $\mathscr{T}$ (for testing). By feeding it all of the pairs $(x, t) \in \mathscr{L}$ of data, we wish to optimize the choice of the network parameters (weights and biases) so that it correctly generalizes in predicting the labels for samples in $\mathscr{T}$. We will do this using standard (stochastig) gradient descent.

The optimization is performed by first introducing a cost function $C(W, b)$ that depends on the weights and biases of the network. It will be defined such that the cost function is minimized when the network correctly predicts the labels. Let us denote the prediction of the network for a given input $\boldsymbol{x}$ by $z(\boldsymbol{x})$, then the usual quadratic cost function is

[2] This really is not a very clever representation, since shifting the whole image by one pixel creates a very different input vector, whereas the original image is still clearly recognizable.

defined via

$$C(W,b) = \sum_{(\boldsymbol{x},t)} \in \mathscr{L} \frac{1}{2} (z(\boldsymbol{x}) - t)^2 . \tag{3.3}$$

Since this cost function is a sum of non-negative terms only, it is clearly minimized when for every input pair we have $z(\boldsymbol{x}) = t$. Notice that the weights and biases are hidden in $z(\boldsymbol{x})$.

We are now in the position to define the backpropagation algorithm, which is nothing but gradient descent. This is most easily imagined if we think of the cost function as having only two parameters, e.g. $C(x_1, x_2)$. The quadratic cost function then defines a 3D landscape of which we wish to find the global minimum. To do so, at any point of this landscape we should move along the direction that points steepest downward (assuming at the moment that there are no local minima). This direction is given by the gradient $\left(\frac{\partial C}{\partial x_1}, \frac{\partial C}{\partial x_2}\right)$, or by ∇C in general. Denoting a given weight between neurons i and j in layers l and $l+1$ respectively as w_{ij}^l, the update rule for our weights is (and similar for the biases):

$$W_{ij}^l \to W_{ij}^l - \alpha \frac{\partial C}{\partial W_{ij}^l}, \tag{3.4}$$

where we have introduced the learning rate α as a hyperparameter. Hyperparameters are parameters that we use as knobs for tweaking the performance of the network. Having a very large α means initially we will make big steps in the direction of the global minimum, but once we are close to it we will most likely overshoot it due to the large stepsize. A very small α will take very long to converge however. In practice, many schemes for choosing and dynamically updating α exist. We will however use a fixed α that we choose based on trial and error.

Notice that the update rule in Eq. 3.4 makes apparent the need for a smooth activation function; it is so that the partial derivative can safely be computed.

EXPLICIT EXAMPLE

For additional clarity, let us explicitly go through an example of the backpropagation algorithm on the simple network in Fig. 3.3. To update the weight w_1, we proceed as

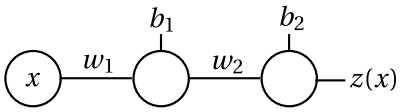

Figure 3.3: A simple network with a single hidden neuron, to demonstrate the backpropagation algorithm. See the main text for the update rule for the weight w_1.

follows. For notational clarity, let us write $z_2 = z(x)$ for the output, and z_1 for the output of the hidden neuron, so that $z_1 = \sigma(x w_1 + b_1)$ and $z_2 = \sigma(z_1 w_2 + b_2)$. The update for w_1

then reads:

$$\begin{aligned} w_1 \to w_1 - \alpha \frac{\partial C}{\partial w_1} &= w_1 - \alpha \frac{\partial \left(\frac{1}{2}(z_2 - t)^2\right)}{\partial w_1} \\ &= w_1 - \alpha \Delta \frac{\partial z_2}{\partial w_1} \\ &= w_1 - \alpha \Delta \frac{\partial \sigma(z_1 w_2 + b_2)}{\partial w_1} \\ &= w_1 - \alpha \Delta z_2' \frac{\partial z_1}{\partial w_1} \\ &= w_1 - \alpha \Delta z_2' z_1' x. \end{aligned} \tag{3.5}$$

In this derivation, we have defined $z'(x) = \sigma'(x)$ as the derivative of $\sigma(x)$ w.r.t. x. The backpropagation algorithm is hence nothing but repeated applications of the chain rule.

Finally, let us discuss two improvements that we can make on the neural network that we have discussed so far. The first is a change in cost function, where instead of the quadratic form we will use the cross-entropy cost function defined as

$$C(W, b) = -\sum_{(\boldsymbol{x}, t)} t \ln z(x) + (1 - t) \ln(1 - z(x)). \tag{3.6}$$

This cost function for the sigmoid neurons has the advantage that the gradient with respect to the weights $\frac{\partial C}{\partial W}$ is proportional to the error:

$$\frac{\partial C(W, b)}{\partial W} \sim \sum_{(\boldsymbol{x}, t)} x \left(z(x) - t\right). \tag{3.7}$$

Hence making a bigger error leads to a bigger correction, which speeds up the process of learning.

The second addition to the network is the use of a regularization parameter l_2. This L2 regularization is added to the cost function as

$$C(W, b) = C_0 + l_2 \sum_w w^2, \tag{3.8}$$

where C_0 is the unregularized cost function and the sum over w symbolically represents summing the elements of the weight matrices. Including this term in the backpropagation algorithm leads to a new update rule for the weights:

$$\begin{aligned} w &\to w - \alpha \frac{\partial C_0}{\partial w} - l_2 w \\ &= (1 - l_2)\, w - \alpha \frac{\partial C_0}{\partial w}. \end{aligned} \tag{3.9}$$

For this reason, this particular type of regularization is also called weight decay. For positive l_2, the network tries to keep the weights small. Intuitively, this means that small

changes in the input also ever only lead to small changes in the output. Effectively therefore, the network becomes insensitive to noise in the input and rather tries to learn patterns.

In the next section, we will make use of neural networks to detect phase transitions based on patterns in the entanglement spectrum.

3.4. CONFUSION SCHEME

Usually, the application of a neural network to a classification problem requires a prelabelled set of data. Here, however, we focus on a combination of supervised and unsupervised techniques. In most cases in physics namely, it is exactly the labelling that one would like to find out (i.e. classification of phases), implying that a labelling is not known beforehand.

In this section we demonstrate that by explicitly mislabelling the data and evaluating the networks' performance, it is possible to identify the correct set of labels. We would stress here, that this method is more involved than an unsupervised approach such as principal component analysis. It is however more powerful in general, with the capability of extracting general non-linear relationships in the data.

To prevent having to deal with exponentially large wavefunctions, we pre-process the wavefunctions to obtain the entanglement spectra [57]. The entanglement spectrum has been argued to contain physical information on the wavefunction, and is hence expected to provide a good representation of it [63]. As we discussed in the previous chapter, its properties can be used to identify symmetry protected topological phases [48, 59, 71]. In fact, it can also provide information on symmetry breaking phases [72, 73] and even many-body localization (a topic that we will discuss in more detail in Chapter 5) [74, 75]. Very recently, an experimental protocol for measuring the ES has been proposed [76]. It is important to remark that various types of bipartition of the whole system into subsets A and B exist, such as dividing the bulk in extensive disconnected parts [58], divisions in momentum space [77] or indeed even random partitioning [61]. We will use the usual spatial bipartition into left and right halves of the whole system.

To justify the use of the ES further, we note that recently quantum entanglement has taken up a major role in the characterization of many-body quantum systems in general [78, 79].

In the following we will introduce a method for discovering phase transitions from unlabelled data with a supervised routine that needs labels. We build this method starting from the Kitaev chain, and then detail it on a classical Ising model and on a disordered interacting spin chain.

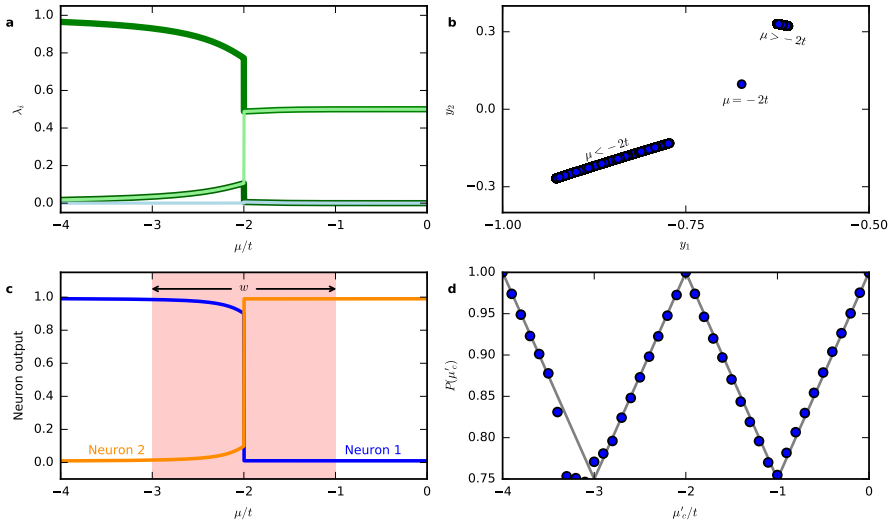

Figure 3.4: Panel **a**: Evolution of the entanglement spectrum as a function of the chemical potential μ. Here we plot the largest four eigenvalues of the reduced density matrix ρ_A. The degeneracy structure is clearly observable. Panel **b**: Principal component analysis of the entanglement spectrum. All data points are shown in the plane of the first two principal components y_1 and y_2. Panel **c**: Supervised learning with blanking. The shaded region is blanked out during the training phase, and the neural network can still predict the correct transition point $\mu = -2t$. Panel **d**: $P(\mu_c')$, evolution of the accuracy of prediction, as a function of the proposed critical point μ_c', which shows the universal W-shape. See text for more details. (Parameters for training: batch size $N_b = 100$, learning rate $\alpha = 0.075$ and regularization $l_2 = 0.001$.)

3.4.1. KITAEV CHAIN

We will develop our method for identifying phase transitions on the example of the Kitaev chain, which serves as an excellent example due to the availability of analutical results. Its Hamiltonian is given by

$$\hat{H} = -t\sum_{i=1}^{L}\left(\hat{c}_{i+1}^{\dagger}\hat{c}_i + \hat{c}_{i+1}\hat{c}_i + h.c.\right) - \mu\sum_{i=1}^{L}\hat{c}_i^{\dagger}\hat{c}_i, \tag{3.10}$$

where $t > 0$ controls the hopping and the pairing of spinless fermions alike and μ is a chemical potential. The groundstate of this model has a quantum phase transition from a topologically trivial ($|\mu| > 2t$) to a nontrivial state ($|\mu| < 2t$) as the chemical potential μ is tuned across $\mu = \pm 2t$.

In panel **a** of Fig. 3.4, we show the first four eigenvalues of the reduced density matrix of an $L = 20$ site Kitaev chain (although we will use the first 10 for the following analysis). The non-trivial phase with even multiplicity degeneracy is clearly distinguishable from the trivial one. This pattern also makes it likely that PCA is able to distinguish the data. The entanglement spectra are interpreted as points in a 10-dimensional space, and we use principal component analysis (PCA) [80] to extract mutually orthogonal axes along

which most of the variance of the data can be observed. PCA amounts to a linear transformation $Y = XW$, where X is an $N \times 10$ matrix containing the entanglement spectra as rows ($N = 10^4$ is the number of samples). The orthogonal matrix W has vectors representing the principal components ω_ℓ as its columns, which are determined through the eigenvalue equation $X^T X\omega_\ell = \lambda_\ell \omega_\ell$. The eigenvalues λ_ℓ are the singular values of the matrix X, and are hence non-negative real numbers, and we normalize them s.t. $\sum \lambda_\ell = 1$. The result of PCA is shown in Fig. 3.4(b), and it is indeed possible to cluster the spectra in to three sets: $\mu < -2t$, $\mu = -2t$, and $\mu > -2t$.

Suffice it to say that the training of a neural network on this data is not necessary at all, but we still do so in order to introduce our proposed confusion scheme. We use an architecture of $[10, 80, 2]$ to learn the trivial and non-trivial phases. The first/second output neuron predicts the (not necessarily normalized) probability for the data to be in trivial/nontrivial phase, and the predicted phase is the phase with the larger probability. The neurons are of the sigmoid-type, with a cross-entropy cross function and we include an L2 regularization term. The network easily learns to distinguish the spectra. To show its generalizing power we blank out the data in a certain width w around the transition point. The data within this region is assigned to the test set $\mathscr{T}$ whereas the network is trained on the data outside this interval, the learn set $\mathscr{L}$. In panel **c** of Fig. 3.4 we show that the network has no difficulty in still prediction the correct transition point even for $w = 3t$.

The PCA as an unsupervised learning technique may be applied without perfectly known information of the system, but it is a linear analysis and is hence incapable of extracting non-linear relationships among the data. On the other hand, a NN is capably of fitting any non-linear function [70], but a training phase with *correctly* labelled input-output pairs is needed. In the following, we propose a scheme combining the two methods which we refer to as a *confusion scheme*.

Suppose all data lies in the parameter range (a, b), and we know there exists a critical point $a < c < b$ such that the data could be classified into two groups. However, we do not know the value of c. We propose a critical point c', and train a network that we call $\mathscr{N}_{c'}$ by labelling all data with parameters smaller than c' with label 0 and the others with label 1. Next, we evaluate the performance of $\mathscr{N}_{c'}$ on the entire data set and refer to its total performance, with respect to the proposed critical point c', as $P(c')$. We will show that the function $P(c')$ has a universal *W-shape*, with the middle peak at the correct critical point c. Applying this to the Kitaev model, we can see from Fig. 3.4d that for $-4t < \mu < 0$, the prediction performance from confusion scheme has a W-shape with the middle peak at $\mu = -2t$.

The W-shape can be understood as follows. We assume that the data has two different structures in the regimes below c and above c, and that the NN is able to find and distinguish them. We refer to these different structures as features. When we set $c' = a$, the NN chooses to assign label 1 to both features and thus correctly predicts 100% of the data. A similar analysis applies to $c' = b$, except that every data point is assigned the label 0. When $c' = c$ is the correct labelling, the NN will choose to assign the right label to both

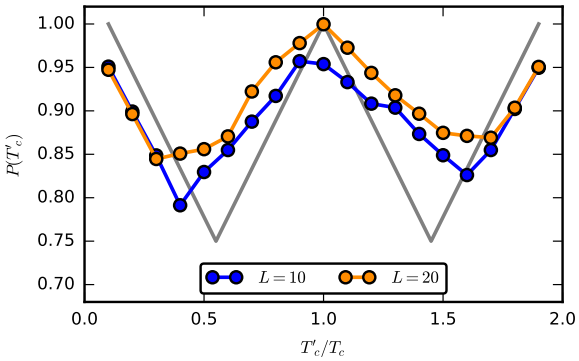

Figure 3.5: Applying the confusion scheme to the Ising model. Position of the middle peak in the universal W-shape deviates from $T'_c = T_c$ for $L = 10$ due to the finite-size effect. Here $k_B T_c \approx 2.27J$ is the exact transition temperature in the thermodynamic limit. For $L = 20$ the middle peak is located exactly at $T'_c = T_c$. (Parameters for training: batch size $N_b = 100$, learning rate $\alpha = 0.02$ and regularization $l_2 = 0.005$.)

sides of the critical point and again performs perfectly. When $a < c' < c$, in the training phase the NN sees data with the *same* feature in the ranges from a to c' and from c' to c, but having *different* labels (hence the confusion). In this case it will choose to learn the label of the majority data, and the performance will be

$$P(c') = 1 - \frac{\min\left(c - c', c' - a\right)}{c - a}. \tag{3.11}$$

Similar analysis applies to $c < c' < b$. This gives the typical W-shape seen in Fig. 3.4d. Notice that if the point c is not exactly centered between a and b, the W-shape will be slightly distorted. Its middle peak always corresponds to the correct labelling, but the depth of the minima will differ between the left and right.

3.4.2. 2D ISING MODEL

We test the confusion scheme on the thermal phase transition in the two-dimensional classical Ising model, which has been studied by both supervised learning [65] and unsupervised learning [66] methods. Here we train a NN (with L^2 neurons in the input and hidden layers, and 2 neurons in the output layer) on the $L \times L$ classical configurations sampled from Monte Carlo simulation. As shown in Fig. 3.5, the W-shape again predicts the right transition temperature. Note the confusion scheme works better when the underlying feature in the data is shaper, i.e. for the larger system size $L = 20$.

To confirm that the confusion scheme indeed extracts nontrivial features from the input data. We have checked the performance curve from confusion scheme, when the NN is trained on unstructured random data. We use a fictive parameter as a tuning parameter, but have completely unstructured (random) data as a function of it. Hence, the network

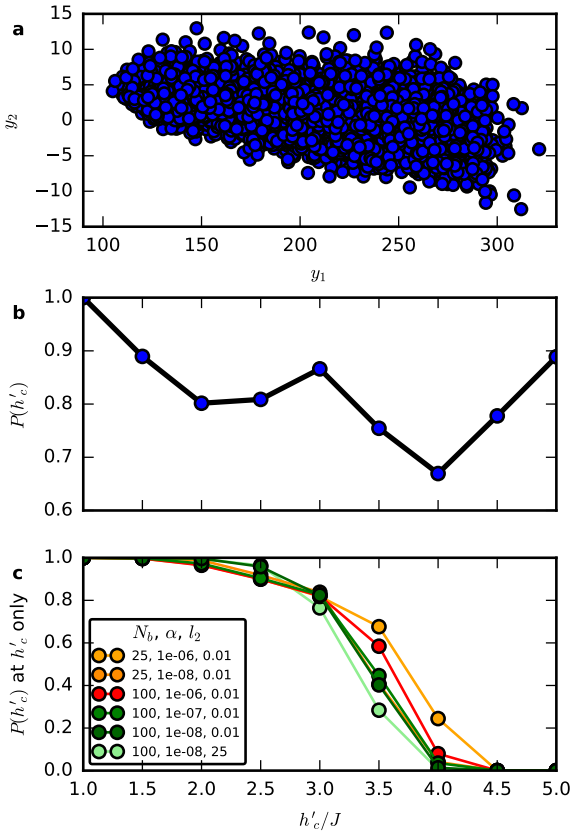

Figure 3.6: Panel **a**: Principal component analysis of the random-field Heisenberg model. Unlike in the Kitaev model or for the Ising data [66], there is no clearly observable clustering. Panel **b**: The characteristic W-shape of the performance curve on the MBL data. The result shows that the network $\mathcal{N}_{h'_c}$ for $h'_c \approx 3J$ performs best, indicating that this is the correct labelling. The distinction between the thermalizing and non-thermalizing phase can hence be put at $h_c \approx 3J$, in agreement with Ref. [75]. (Parameters for training: batch size $N_b = 100$, learning rate $\alpha = 10^{-8}$ and regularization $l_2 = 0.01$.) Panel **c**: The performance of network $\mathcal{N}_{h'_c}$, when evaluated at the point h'_c only instead of on the full data, for various different sets of learning parameters (see legend). Clearly the performance of the network is most independent of the exact training scheme at $h'_c \approx 3J$, showing a robustness of this correct labelling against variations in training.

will not find structure in the data, and a correct labelling does not exist. The middle peak of the characteristic W-shape disappears, turning it into the V-shape shown in Fig. 3.7.

We notice that the choice of the learning rate (α) and regularization (l_2) is essential for a successful training. As mentioned before, the use of regularization is expected to reduce overfitting and make the network less sensitive to small variations of the data, hence forcing it to rather learn its structure [81]. However, the confusion scheme depends

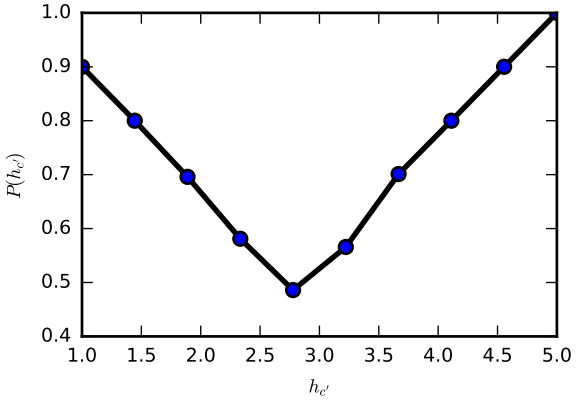

Figure 3.7: Applying the confusion scheme to random data turns the characteristic peaked W shape into a V shape.

solely on the ability of finding the majority label for the underlying structure in the data. In this sense, overfitting is not necessarily bad. Indeed we have observed that training with a negative l_2 may lead to an equally good performance. We speculate that this is because a negative l_2 tries to quickly increase the weights, making it harder for the network to change its opinion about data samples in later stages. If the initial training data is uniformly sampled, meaning the majority data is indeed represented by a majority, the network will rapidly adjust its weights to this majority. Lastly, we mention that trainings are performed in epochs. In each epoch all training data is passed once in batches of size N_b in a random order. The training is stopped when a clear W-shape is formed.

3.4.3. RANDOM-FIELD HEISENBERG CHAIN

We will now test our proposed scheme on an example where we the exact location of the transition point is now known. We study a case of interest in recent literature, namely that of many-body localization. We consider the following model:

$$H = J \sum_{i=1}^{L} \boldsymbol{S}_i \cdot \boldsymbol{S}_{i+1} + \sum_{\alpha=x,y,z} \sum_{i=1}^{L} h_i^\alpha S_i^\alpha, \tag{3.12}$$

where $\boldsymbol{S}$ denote spin-1/2 operators. The local fields h_i^α are drawn from a uniform box distribution with zero mean and width h_{max}^α. We set $h_{\text{max}}^x = h_{\text{max}}^z = h_{\text{max}}$ and $h_{\text{max}}^y = 0$. The disorder allows us to generate many samples at a fixed set of model parameters, in analogy to the different configurations for a fixed temperature in the classical spin systems [65, 66].

The subject of many-body localization will be explained in much greater detail in Chapter 5. For our current discussion, it suffices to highlight that the model in equation

(3.12) has a transition driven by the disorder strength $h_{\max}$. In particular, when varying $h_{\max}$, both the energy level statistics as well as the statistics of the entanglement spectra change their nature [75]. For the case of the energy levels, the gaps (level spacings) follow either a Wigner-Dyson distribution for the thermalizing phase, or a Poisson distribution for the localized phase (we refer to sections 5.2.2 and 5.2.2 for more detail). For the entanglement spectrum, the Wigner-Dyson distribution is replaced by a semi-Poisson distribution. Note that the change of ES can already be seen from the statistics in a *single* eigenstate [75].

We numerically obtain the entanglement spectrum for the groundstate of the model in equation (3.12), for disorder strengths between $h_{\max} = J$ and $h_{\max} = 5J$. The transition was shown to happen around $h_{\max} \approx 3J$ [75], but we stress that our method does not rely on this knowledge. We would simply have started from a larger width of points, and then systematically narrow it down to the current range. At each value of $h_{\max}$ we generate 10^5 disorder realizations for system size $L = 12$ and calculate the entanglement spectrum for $L_A = L_B = 6$. These $2^6 = 64$ levels are used as the input to the NN.

First though, we try to use an unsupervised PCA to cluster the data. This analysis shows that the first two principal components are dominant, with the other components being of order 10^{-4} or less. A scatterplot of the data when projected onto the first two principal components (shown in panel **a** of Fig. 3.6) does not reveal a clear clustering of the spectra however.

We therefore turn to train a shallow feedforward network on the entanglement spectra in order to use the confusion scheme. Here we use a network with 64 input neurons, 100 hidden neurons and 2 output neurons. The results are shown in Fig. 3.6**b**. Also in this case, the characteristic W-shape is obtained and we detect the transition at $h_c \approx 3J$. In addition to the previous cases, we also consider explicitly the performance of the network $\mathcal{N}_{h_c'}$ *at* h_c'. We do this to confirm that the labelling with h_c' at $3J$ is indeed correct. We expect namely that the training of the network is most robust against changes in its parameters for the correct labelling. In other words, we may also look for the h_c' at which the training is most independent of chosen conditions. As shown in Fig. 3.6**c**, this point is also at $h_c' \approx 3J$.

3.4.4. DISCUSSION

We have explored the uses of machine learning in identifying (quantum) phase transitions from the overall performance of neural networks. As input data for training we have used the entanglement spectrum of a quantum state, since it provides an excellent way of compressing the otherwise exponentially large wavefunction. We have shown that by confusing the neural network, i.e. by purposefully mislabelling the data, we are able to identify the phase transition between two phases of matter without prior knowledge of them. We demonstrated this for three different scenario's. We expect our method to be useful also in situations where it is not clear whether two phases are identical, or whether intervening phases exist in between other phases. If the underlying data in such situa-

tions has a (hidden) pattern that distinguishes them, our machine learning approach may well be able to find it. Our method is not limited to applications in condensed matter physics and phase transitions, but might prove useful in any field in which large amounts of data can be classified using a priori unknown labels.

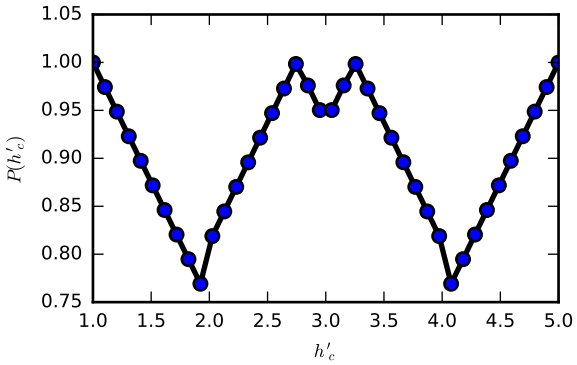

Figure 3.8: As a function of a parameter $h \in [1,5]$, we create data that has different underlying structures in the regimes $1 < h \leq 2.75$, $2.75 < h \leq 3.25$ and $3.25 < h \leq 5$. From the peak strcture obtained from the performance of the networks, we are able to identify each of these three regimes.

An interesting direction for future studies is the relaxation of the assumption that there are only two phases to be distinguished. In fact, we may always split the data into two sets A and B and ask the neural network to distinghuish 'in set A' or 'in set B'. We verify this by artificially creating data with three underlying structures as a function of a fictional parameter. The characteristic W-shape is changed into a structure that has two central peaks, indicating the two phase transitions. We demonstrate this in Fig. 3.8. It may be possible to formulate this method in a self-consistent way, with an adaptive labelling and having the algorithm determine the correct labels by itself.

III

OUT-OF-EQUILIBRIUM

4

MIXED STATE TOPOLOGY

In this chapter we move from pure states to mixed states, and consider the application of the projective symmetry representation analysis on density matrices. To that end we decompose the density matrix into its amplitudes and investigate the action of the symmetry operators on those. Moving to a matrix product state representation for mixed states also allows us to study the effect of environmental noise on symmetry protected topological states, and we demonstrate that for Markovian baths coupling to every site of the system the topological nature is lost.

4.1. GENERAL

Parts of this chapter have been published as Ref. [82]. We study the extension of the classification of symmetry protected topological groundstates in closed systems to that of open systems. In addition to the classification, we study the loss of topological nature due to the coupling of the system to different types of environments. The numerical tools that were developed for the purpose of this work are shortly described, but its details are defered to Appendix A. The required background for the projective symmetry group analysis is covered in Chapter 2 and Appendix B.

4.1.1. MOTIVATION

The classification of groundstates that are separated from the lowest excited states by an energy gap has received a considerable amount of (renewed) interest during the last decade. This was due to the realisation that not all such states can be classified using the Ginzburg-Landau scheme. In particular, there are groundstates that are invariant under the same set of symmetries, but which can nevertheless not be deformed into eachother without closing the energy gap. These states are then called topologically distinct. Prime examples are the fractional quantum Hall states [83] or the ground states of certain spin systems [84].

The topologically non-trivial states of matter often have remarkable properties such as non-abelian quasi-particles and robust edge modes. This makes this state of matter one that is sought after intensely, both in classical materials as well as in quantum engineered systems like cold atoms and superconducting qubits. Especially in such engineered systems, the high level of control they achieve often comes at the price of coupling the system to the environment. This raises the question of whether topologically non-trivial states are stable to such a coupling, and moreover it raises the question of whether it is possible to classify non-equilibrium topological states at all [85–87]. In this chapter we adress these questions. We study the evolution of topological groundstates in the presence of non-equilibrium noise and propose a classification scheme for the corresponding mixed density matrices in one spatial dimension.

Before we turn to the non-equilibrium setting, let us quickly mention that the identification of topological groundstates at zero temperature in a closed system is well understood [83, 88, 89]. The method of identifying the topological properties depends on the system at hand: For the case of insulating free fermions one turns to topological invariants [9, 10, 89]. For interacting systems general features such as edge states, ground state degeneracies and excitation statistics are indicators of topology [90]. In 1D however, since the only non-trivial phases in 1D are symmetry protected [91], we turn to the methods we have discussed in Chapter 2. In other words, we will make use of projective representations of symmetry groups; this provides a powerful tool for classifying 1D topological phases [42, 49].

Continuing in the same spirit as of Chapter 2, we also discuss the entanglement (spec-

trum) properties and its relation to topological states. For classical systems, correlation functions can be used to identify phases and phase transitions. It is only natural to expect that "quantum correlations", i.e. entanglement, shows features of quantum phase transitions.

We will extend both the projective symmetry group analysis and the entanglement spectrum to the case of mixed states. For the latter, we identify a new type of topologically non-trivial mixed state based on the degeneracies. We will then use both of these methods to assess the stability of a topological phase in an open system. A number of attempts exist in extending the topological invariants to open systems [92–94]. These methods are restricted to specific types of master equations, however, and do not provide a general tool for the identification of open system topology.

For concreteness, we develop and demonstrate our ideas on the example of the antiferromagnetic spin-1 chain of Eq. 2.14, reproduced here for convenience:

$$H = J\sum_i \mathbf{S}_i \cdot \mathbf{S}_{i+1} - B\sum_i S_i^x - D\sum_i \left(S_i^z\right)^2. \tag{4.1}$$

The $\mathbf{S}$ operators represent $S = 1$ degrees of freedom, B sets the strength of an external field and D is an on-site anisotropy. We measure both B and D in units of $J = 1$. For an extended region in the B–D phase diagram (shown in Fig. 2.3) around the Heisenberg point ($B = D = 0$), this model is in a symmetry-protected topological phase called the Haldane phase [41].

The symmetries responsible for the protection of Haldane phase are time reversal and inversion symmetry [41]. Along the $B = 0$ line, the model also has a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry discussed in section 2.2.3, allowing for the definition of a non-local order parameter called string order [44] (c.f. Eq. 2.13):

$$\langle O^{x,z}\rangle = \lim_{|i-j|\to\infty} \langle\psi|\mathbf{S}_i^{x,z} e^{-i\pi\sum_{i<k<j}\mathbf{S}_k^{x,z}} \mathbf{S}_j^{x,z}|\psi\rangle. \tag{4.2}$$

The string order parameter in z detects a dilute antiferromagnetic ordering of the spins in the z-direction, where the quantum numbers m^z alternate perfectly between $+1$ and -1 after stripping out the $m^z = 0$ values (and similarly for the string order in x). For this model, a non-vanishing string order identifies the Haldane phase. For the cases of other symmetries than $\mathbb{Z}_2 \times \mathbb{Z}_2$, a generalized string order can be defined [46]. Unfortunately, string order is restricted to 1D systems and cannot be used to fully *classify* 1D phases [42].

The numerical approach that we will be using for the rest of this chapter is based on matrix product states. The MPS formalism gives us easy access to the projective symmetry group analysis, and in addition the entanglement spectrum $\{\lambda_\alpha\}$ is obtained for free. We have seen in section 2.5 that it can be used as a tool for studying topoplogical phases. In the 1D SPT phases, the feature in the ES of a non-trivial state is that all the values λ_α form pairs of even multiplicity [47]. This is fundamentally due to the projective symmetry representation matrices U_g being antisymmetric [47]. However, a degenerate ES is not in one to one correspondence with the topology of the state. The ES is hence easily

available but not a strictly conclusive tool for investigating gapped symmetric phases. For example, a state symmetric under both $\mathbb{Z}_2 \times \mathbb{Z}_2$ and inversion could be trivial (in the sense of projective symmetry representations) with respect to one but not the other; both would induce the same degeneracy in the ES.

4.2. METHODS

The addition of an environment in the description of the system means that we must use the language of density matrices rather than pure states. Fortunately, the formalism of MPS and time evolution using the (i)TEBD algorithm may be extended to the case of mixed states. The details on this extension can be found in Appendix A, but we review the necessary concepts here. The evolution of a density matrix ρ is governed by a master equation in Lindblad [20] form:

$$\dot{\rho}(t) = \mathscr{L}[\rho] = -\frac{i}{\hbar}[H,\rho] + \gamma \sum_i \left[L_i^\dagger L_i \rho - \frac{1}{2}\left\{\rho, L_i L_i^\dagger\right\}\right], \tag{4.3}$$

in which the jump operator L represents the system-part of the coupling to the environment. We shall consider different jump operators in our discussion of the results. We first turn to the simulation of the density matrix.

The MPS formalism can be extended to simulating the density matrix via the superoperator approach [95, 96]. To this end, we introduce a set of local basis matrices denoted σ_i. These matrices are $d \times d$ matrices, with d being the dimension of the local Hilbert space. A convenient (and in fact, numerically most efficient) choice for these basis matrices is the set of generators for $SU(d)$, augmented by the identity matrix. This is because the generators are traceless, leading to only certain terms contributing to expectation values (see Appendix A). The density matrix is then expanded in this basis, and the coefficients are written in MPS form as:

$$\rho = \sum_{ijk\ldots} \rho_{ijk\ldots} \sigma_i \sigma_j \sigma_k \cdots, \tag{4.4}$$

$$\rho_{ijk\ldots} = \mathrm{Tr}\left(M_i^{[1]\sharp} M_j^{[2]\sharp} M_k^{[3]\sharp} \cdots \right). \tag{4.5}$$

The density matrix ρ is then interpreted as a vector $|\rho\rangle_\sharp$, on which superoperators act as matrices. A superoperator T_A that represents multiplication by a matrix A, i.e. $T_A[\rho] = A\rho$, is turned into a matrix $T_A^\sharp$ acting on the vectorized density matrix as[1]

$$T_A^\sharp |\rho\rangle_\sharp = |T_A[\rho]\rangle_\sharp = |A\rho\rangle_\sharp. \tag{4.6}$$

Using this formalism for the master equation in Eq. 4.3, it turns into $|\dot{\rho}\rangle_\sharp = \mathscr{L}_\sharp |\rho\rangle_\sharp$. This equation for the time evolution of the vector $|\rho\rangle_\sharp$ is formally equivalent to the Schrödinger

[1] It is important to note here that the resulting state of such an operation is not necessarily a positive semi-definite matrix. Practically therefore, one has to continuously check that the state under consideration is still a proper density matrix.

equation for the time evolution of a vector $|\psi\rangle$. This allows us to simulate the density matrix MPS using the same algorithms as for the pure state MPS [50, 95, 97].

There are a few important differences to the pure state case however. Most notably, from the MPS matrices of the pure state we are able to compute the entanglement spectrum. It is not obvious what the corresponding object is when we perform the same routine on the matrices for the density matrix. This set of values on the density matrix level is referred to as the operator space entanglement spectrum (OSES), denoted $\lambda_\sharp$. It is obtained from the $A^{[n]\sharp}$ matrices via a decomposition into $\Gamma_\sharp \lambda_\sharp^{[n]}$ in the same way as the ES is obtained in the pure state case.

Next, let us analyze what we can learn from the OSES. If $|\rho\rangle_\sharp$ describes a pure state, it is straightforward to show that $\lambda_\sharp = \lambda \otimes \lambda$ (see Eq. A.40). For a general mixed state the relation between the two becomes more involved. Fortunately, the degeneracy (pairs of even mutliplicity) of the ES is a property that comes from a symmetry of the state, and hence persists also for the OSES. This allows us to directly interpret the degeneracies of the OSES and identify topologically non-trivial states. In the following sections, we demonstrate the use of the superoperator approach on the spin-1 Heisenberg model of Eq. 4.1.

4.3. RESULTS

To obtain the groundstate density matrix for a given Hamiltonian, we perform a 'cooling' procedure starting from the infinite temperature state. This is based on the ability of calculating any thermal density matrix using the superoperator approach [96]. The infinite temperature state is one of the easiest states to represent using the MPS formalism, since ρ is just a normalized identity matrix. In the expansion in the basis states σ_i that means on every site we only need the identity matrix, which we choose to be the σ_0 matrix. We hence need only one non-zero coefficient, namely $\rho_{00\ldots0}$ in Eq. 4.5.

We denote the infinite temperature state $|1\rangle_\sharp$. Using this state as a starting point, we can reach any finite temperature state as follows. We introduce the superoperator $T[A] =$

$\frac{1}{2}(AH+HA)$ with H the Hamiltonian. Then consider the calculation:

$$\begin{aligned} e^{-\beta T_\sharp}|1\rangle_\sharp &= \left(1-\beta T_\sharp + \frac{1}{2}\beta^2 T_\sharp^2 + \ldots\right)|1\rangle_\sharp \\ &= |1\rangle_\sharp - \beta T_\sharp|1\rangle_\sharp + \frac{1}{2}\beta^2 T_\sharp^2|1\rangle_\sharp + \ldots \\ &= |1\rangle_\sharp - \beta|T[1]\rangle_\sharp + \frac{1}{2}\beta^2 T_\sharp|T[1]\rangle_\sharp + \ldots \\ &= |1\rangle_\sharp - \beta|H\rangle_\sharp + \frac{1}{2}\beta^2|T[H]\rangle_\sharp + \ldots \\ &= |1\rangle_\sharp - \beta|H\rangle_\sharp + \frac{1}{2}\beta^2|H^2\rangle_\sharp + \ldots \\ &= |e^{-\beta H}\rangle_\sharp \end{aligned} \tag{4.7}$$

Hence, performing 'time'-evolution up to a time β, using the superoperator $T[A]$ after starting from the infinite temperature state, gives us any thermal state at temperature β.

The results for the approach of the energy (red curve) towards the ground state at the Heisenberg point ($B = D = 0$) are presented in the top panel of Fig. 4.1, as a function of the inverse temperature β. The figure also shows the emergence of the string order parameter $\langle O^{x,z}\rangle$. Additionally, it shows $\mathcal{O}_{Z_2\times Z_2}$, computed from the represenation matrices U_g of the $\mathbb{Z}_2\times\mathbb{Z}_2$ symmetry on the MPS as:

$$\mathcal{O}_{Z_2\times Z_2} = \frac{1}{\chi}\mathrm{Tr}(U_x U_z U_x^\dagger U_z^\dagger) = \begin{cases} 0 \\ \pm 1 \end{cases}, \tag{4.8}$$

That is, $\mathcal{O}_{Z_2\times Z_2}$ takes the value 0 if the state is not symmetric, or ± 1 for a symmetric and trivial or non-trivial state, respectively. These measurements were performed using an MPS of maximal bond dimension $\chi = 100$, keeping the truncation error below 10^{-8} at all times. The emergence of non-zero string order in the top panel of Fig. 4.1 happens at non-zero temperature. Even though the Mermin-Wagner theorem states that no order at finite temperature exists in 1D, it is not immediately obvious that this holds for non-local order. he bottom panel of the figure shows the temperature β_c at which string order develops as a function of inverse system size. By performing this finite size scaling, we may conclude that indeed in the thermodynamic limit the string order only appears at $T = 0$.

4.3.1. PROJECTIVE SYMMETRY GROUP ANALYSIS

There is an important difference between the evaluation of the projective symmetries on the MPS of a pure state to that of a mixed state.

To illustrate why the direct approach of Chapter 2 fails, let us decompose the density matrix into its 'amplitudes':

$$\rho = \omega^\dagger \omega.$$

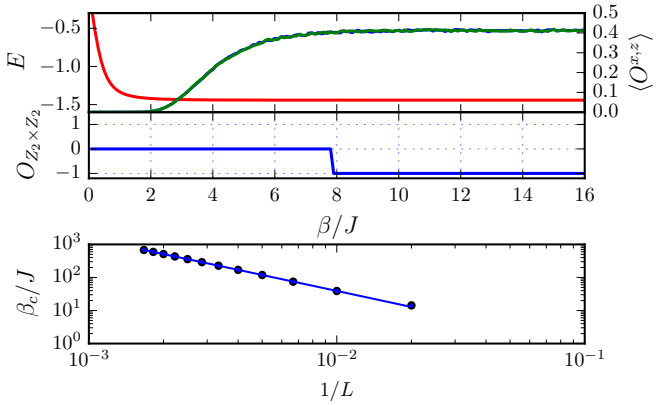

Figure 4.1: The top panel shows the convergence of energy (red) and the increase of string order (blue and green) while cooling towards the groundstate. The $\mathcal{O}_{Z_2 \times Z_2}$ trace shows the state going from non-symmetric to symmetric and topologically non-trivial. The bottom panel shows that for finite systems of size L there is no finite temperature β_c at which the string order correlation length diverges, i.e., string order only develops at $T = 0$.

This choice for ω is not unique, since $U\omega$ with U a unitary matrix is a valid amplitude for the same density matrix too. There is one obvious choice for ω though, which is simply $\omega = \sqrt{\rho}$. The natural extension of the projective symmetry group analysis is to write the action of a symmetry operation $g \in \mathcal{G}$, represented as R_g, via

$$\rho \to \tilde{\rho} = R_g \rho R_g^\dagger = \left(R_g \omega\right)\left(R_g \omega\right)^\dagger .$$

As we will see shortly, this approach fails in classifying symmetry protected topological states, since we will always find a trivial factor set for the representation on the MPS level due to the 'double' application of R_g and its inverse $R_g^\dagger$.

Let us first compute the MPS representation of a density matrix $\rho = |\psi\rangle\langle\psi|$ given the MPS representation of $|\psi\rangle$. Letting s_i enumerate the local basis states as usual, we denote the MPS for $|\psi\rangle$ as

$$|\psi\rangle = \sum_{s_1 \ldots s_L} \mathrm{Tr}\left(M^{[1],s_1} \cdots M^{[L],s_L}\right) |s_1 \ldots s_L\rangle .$$

Then a straightforward computation gives us

$$\begin{aligned}
\rho = |\psi\rangle\langle\psi| &= \sum_{\substack{s_1\dots s_L\\ r_1\dots r_L}} \mathrm{Tr}\left(M^{[1],s_1}\cdots M^{[L],s_L}\right)\mathrm{Tr}\left(M^{[1],r_1}\cdots M^{[L],r_L}\right)^*|s_1\dots s_L\rangle\langle r_1\dots r_L| \\
&= \sum_{\substack{s_1\dots s_L\\ r_1\dots r_L}} \mathrm{Tr}\left(M^{[1],s_1}\cdots M^{[L],s_L}\right)\mathrm{Tr}\left(\left(M^{[1],r_1}\right)^*\cdots\left(M^{[L],r_L}\right)^*\right)|s_1\dots s_L\rangle\langle r_1\dots r_L| \\
&= \sum_{\substack{s_1\dots s_L\\ r_1\dots r_L}} \mathrm{Tr}\left(M^{[1],s_1}\cdots M^{[L],s_L}\otimes\left(M^{[1],r_1}\right)^*\cdots\left(M^{[L],r_L}\right)^*\right)|s_1\dots s_L\rangle\langle r_1\dots r_L| \\
&= \sum_{\substack{s_1\dots s_L\\ r_1\dots r_L}} \mathrm{Tr}\left(M^{[1],s_1}\otimes\left(M^{[1],r_1}\right)^*\cdots M^{[L],s_L}\otimes\left(M^{[L],r_L}\right)^*\right)|s_1\dots s_L\rangle\langle r_1\dots r_L| \\
&= \sum_{\substack{s_1\dots s_L\\ r_1\dots r_L}} \mathrm{Tr}\left(N^{[1],s_1,r_1}\cdots N^{[L],s_L,r_L}\right)|s_1\rangle\langle r_1|\otimes\dots\otimes|s_L\rangle\langle r_L|. \qquad (4.9)
\end{aligned}$$

Here we have made use of the relation $\mathrm{Tr}(A\otimes B) = \mathrm{Tr}A\mathrm{Tr}B$, and have introduced a new set of matrices

$$N^{[i],s_i,r_i} = M^{[L],s_L}\otimes\left(M^{[L],r_L}\right)^*,$$

which are of size $D^2\times D^2$ if the original M are of size $D\times D$. Rather than keeping the s and r indices explicitly, we will use the form of Eq. 4.5, keeping in mind that for the pure state case the MPS matrices are obtainable from a direct product.

For the projective symmetry group analysis, we now ask ourselves how the matrices $N^{[i]}$ transform under the application of a local unitary symmetry. From the transformation rule of basis states in Eq. 2.18, we find:

$$|s_i'\rangle\langle r_j'| \to \sum_{s_i,r_j}\left(R_g\right)_{s_i,s_i'}|s_i\rangle\langle r_j|\left(R_g\right)^\dagger_{r_j,r_j'}. \qquad (4.10)$$

Which, in complete analogy to the derivation in section 2.4 leads to the transformation rule of the MPS matrices:

$$\tilde{N}^{[i],s_i,r_i} = e^{i\phi_g}W_g^{-1}N^{[i],s_i,r_i}W_g,$$

with W_g the projective representation. Since we also know how the symmetries act on the $M^{[i]}$ matrices, we may write the transformation as

$$\begin{aligned}
N^{[i],s_i,r_i}\to\tilde{N}^{[i],s_i,r_i} &= \tilde{M}^{[i],s_i}\otimes\left(\tilde{M}^{[j],r_j}\right)^* \\
&= \left(e^{i\theta_g}U_g^{-1}M^{[i],s_i}U_g\right)\otimes\left(e^{-i\theta_g}\left(U_g^{-1}\right)^*\left(M^{[j],s_j}\right)^*\left(U_g\right)^*\right) \\
&= U_g^{-1}\otimes\left(U_g^{-1}\right)^*\ N^{[i],s_i,r_i}\left(U_g\otimes U_g^*\right). \qquad (4.11)
\end{aligned}$$

Comparing the above two results, we define $W_g\equiv U_g\otimes U_g^*$. This still defines a projective representation, as we can deduce by considering the action of two symmetry operations g and h sequentially:

$$\tilde{N}^{[i]} = W_g^{-1}W_h^{-1}N^{[i]}W_hW_g, \quad \text{versus} \quad \tilde{N}^{[i]} = W_{gh}^{-1}N^{[i]}W_{gh}.$$

Hence we may have

$$W_g W_h = e^{i\Phi(g,h)} W_{gh},$$

showing that indeed the W in general form a projective representation. The same is true for the U matrices, i.e. $U_g U_h = e^{i\phi(g,h)} U_g h$. Let us explicitly compute the above two expressions by substituting $W_g = U_g \otimes U_g^*$:

$$\begin{aligned} W_g W_h &= \left(U_g \otimes U_g^*\right)\left(U_h \otimes U_h^*\right) \\ &= U_g U_h \otimes U_g^* U_h^* \\ &= U_g U_h \otimes \left(U_g U_h\right)^* \\ &= e^{i\phi(g,h)} U_g h \otimes e^{-i\phi(g,h)} U_{gh}^* \\ &= \left(U_{gh} \otimes U_{gh}^*\right). \end{aligned} \tag{4.12}$$

This has to be compared to

$$W_g W_h = e^{i\Phi(g,h)} W_{gh} = e^{i\Phi(g,h)} \left(U_{gh} \otimes U_{gh}^*\right), \tag{4.13}$$

from which we can directly see that $\Phi(g,h) = 0 \mod 2\pi$, i.e. the factor set is always trivial. This means that even though the pure state $|\psi\rangle$ may have a non-trivial factor set, this information is lost upon going to a density matrix description. We must therefore turn to a slightly modified approach of classifying topology from density matrices.

Motivated by the cancellation of the phases in Eq. 4.12, we propose to apply the symmetry operator to ρ only on one side, i.e.

$$\rho \to \tilde{\rho} = R_g \rho = R_g \omega^\dagger \omega.$$

This has the additional advantage that the numerical procedure remains unchanged, i.e. it is as if we are applying the symmetry operator to the physical legs of the density matrix MPS. Following the same derivation as in the previous case, we obtain $W_g = U_g \otimes 1$ with 1 representing the identity matrix. This prevents the phase cancellation due to the U_g^* factor from before and means that on the one hand $W_g W_h = e^{i\Phi(g,h)} W_{gh}$ and on the other

$$\begin{aligned} W_g W_h &= \left(U_g \otimes 1\right)\left(U_g \otimes 1\right) \\ &= U_g U_h \otimes 1 \\ &= e^{i\phi(g,h)} U_{gh} \otimes 1. \end{aligned} \tag{4.14}$$

We may conclude that $\Phi(g,h) = \phi(g,h)$, so that the factor set carries over directly from the wavefunction MPS. In this limit, where $\rho = |\psi\rangle\langle\psi|$, we hence recover the correct factor set. This approach is not limited to that case however, and the numerical evaluation of this procedure seems to correctly correlate with the degeneracies of the (operator space) entanglement spectrum.

degeneracy ($n \in \mathbb{N}$, $n > 0$)		significance	color
1	•○○○○○○○○○…	trivial	red
$2n-1$	○○•○•○•○•○…	trivial	yellow
$2(2n-1)$	○•○○○•○○○•…	non-trivial/mixed	blue
$4n$	○○○•○○○•○○…	non-trivial/pure	green

Table 4.1: Color code for OSES plots.

4.3.2. OPERATOR SPACE ENTANGLEMENT SPECTRUM

Let us now turn back to the discussion of the OSES. From our discussion on the ES in Chapter 2, we know that the groundstate of the Heisenberg point in the Haldane phase has at least a two fold degeneracy. For the OSES, since the groundstate is a pure state, this means that the degeneracy should be at least four, since $\lambda_\sharp = \lambda \otimes \lambda$ (see Eq. A.40). In the left panel of Fig. 4.2 the OSES is indeed seen to develop this fourfold multiplicity.

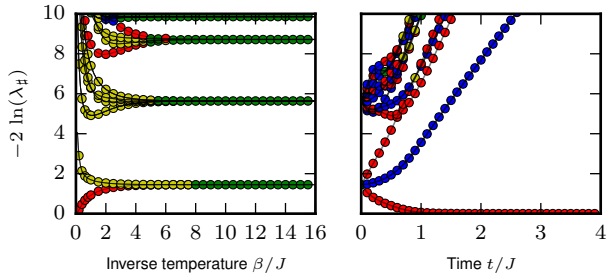

Figure 4.2: (**Left**) The OSES as a function of inverse temperature (only the 42 lowest values are shown). When the temperature is lowered, the expected pairs with multiplicities being multiples of four develops. For the color code, see Tab. 4.1 (**Right**) Time evolution of the ground state at the Heisenberg point ($B = D = 0$) with noise strength $\gamma = 0.5$ for a local jump operator S^- pumping the spins into the minimal S^z eigenstate. Only the lowest 20 values are shown.

4.3.3. TIME EVOLUTION

After obtaining the groundstate of the Heisenberg point by cooling from the infinite temperature state, we will now simulate the master equation of Eq. 4.3 with various jump operators L. The choice of the jump operators is governed by two considerations: Is it physically relevant, and does it break the protecting $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry?

We will consider three different jump operators, the first of which serves as a benchmark. Namely, we consider a jump operator of the form

$$L_i = S_i^- . \tag{4.15}$$

Since we are not driving the system, the steady state of the system under such an evolution can easily be predicted: all of the spins relax to the minimum $m_z = -1$ value. The right panel of Fig. 4.2 shows the time evolution of the OSES in this scenario, starting at $t = 0$ from the Heisenberg groundstate. We observe an immediate splitting of the degeneracies of the levels, and in particular observe only a single value moving towards 0. In other words the OSES evolves to have $\lambda_\sharp = (1, 0, 0, \ldots)$ corresponding to a product state, fully compatible with our expectations that the steady state is a trivial state with all spins polarized. Moreover, the immediate splitting of the degeneracy suggests a breaking of the symmetry protecting the topological degeneracy.

As the next jump operator, we consider the influence of a fluctuating external field,

$$L_i = S_i^z. \tag{4.16}$$

This jump operator also breaks the protecting $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, which is reflected in the trace of $\mathcal{O}_{\mathbb{Z}_2\times\mathbb{Z}_2}$ shown at the bottom of the left panel in Fig. 4.3. The top part of this figure shows the time-evolution of the OSES however, which shows an interesting pattern of degeneracies. Namely, for short times the spectrum shows a set of full even-multiplicity degeneracies but with some pairs being two-fold degenerate. This is incompatible with this spectrum being of the form $\lambda \otimes \lambda$, but nevertheless satisfies all the criteria for being topologically nontrivial. This is therefore a genuinely mixed state with topological OSES that has no pure counterpart. Unfortunately, at longer times the topological nature is completely lost, when the twofold degenerate lowest level splits into two nondegenerate levels.

Another physically relevant and potentially important term is a fluctuating Ising coupling

$$L_i = S_i^z S_{i+1}^z. \tag{4.17}$$

For an implementation of a spin-chain with trapped ions this correspond to fluctuations in the control gates responsible for the couplings [98]. This type of jump operator preserves the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, which is reflected in the initial period where $\mathcal{O}_{\mathbb{Z}_2\times\mathbb{Z}_2}$ remains at -1. The point where this parameter jumps to the value 0, indicating the loss of the symmetry, coincides with the point where the OSES loses its evenfold degeneracies. This is shown in the right panel of Fig. 4.3. In all of the above, the coupling strength was fixed to $\gamma = 0.5$. Changing this value only has an influence on the timescales, and hence may be used to stretch the transient period. However, the type of Lindblad equation that we have used with a single type of jump operator and no driving, basically corresponds to a coupling to an infinite temperature bath. In other words, the system is allowed to exchange energy on all scales. This makes it unsurprising that the topological nature of the system is lost under such an evolution. In fact, when comparing the panels of Fig. 4.3 we may recognize a similar pattern after the transient regime has passed, indicating that the steady state that would be obtained after even longer simulation times is most likely a thermal state.

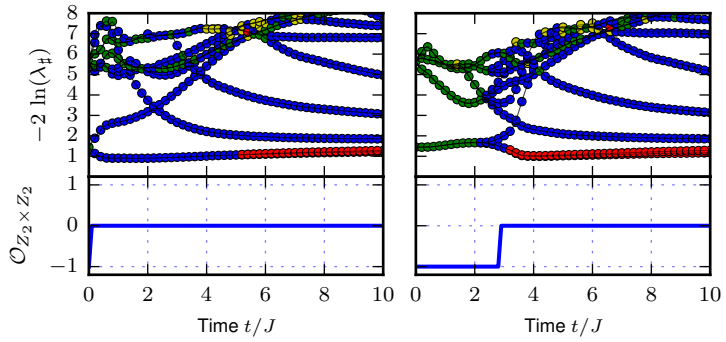

Figure 4.3: Operator space entanglement spectra for $L = S^z$ (left) and $L = S^z S^z$ (right), starting from $B = D = 0$ with $\gamma = 0.5$. Only the lowest 20 values are shown.

4.4. CONCLUSION AND DISCUSSION

Since no physical system is completely isolated, the study of non-equilibrium effects due to the environment is an important issue. In this chapter, we have addressed this question in the setting of topologically non-trivial states in 1D. We have found that we can extend the classification by projective symmetry representations to that of mixed states, justifying the use of degeneracies of the operator space entanglement spectrum as an indicator for topology.

In addition, by coupling the non-trivial groundstate of the Heisenberg model to an environment via suitable types of noise, the mixed state may retain a notion of the topological state in the intermediate regime. This is a clear indication of topological mixed states. On longer timescales, the topological nature is always lost under the type of Lindblad equation we considered. This is most likely due to the availability of processes involving the bath, that may overcome the energy gap protecting the groundstate. The bath coupling strength γ hides all of the structure of the bath and simply sets the timescale for reaching steady state. But having this information compressed to a single number, makes for the bath we consider to effectively represent an infinite temperature reservoir. Hence, on long enough timescales there are always processes that may overcome the energy gap and wash out the topological nature of the state.

5

MANY-BODY LOCALIZATION AND DEPHASING

This chapter focuses on thermalization of quantum systems, or indeed the absence of it. Preventing a quantum system from thermalizing has potential applications as data storage, since information contained in the initial state may be preserved. We discuss what it means for a closed system to thermalize, and what it means for a closed system to many-body localize. In the non-equilibrium spirit, we then investigate the decoherence of a single spin coupled to a bath. In this case, the bath will be a system that has both a thermalizing and many-body localizing regime. These two different regimes can be clearly distinguished in the decoherence properties of the probe spin.

5.1. GENERAL

With the increase in the ability to precisely control and manipulate atoms and ions in the laboratry, there has been a surge in the research on fundamental questions in quantum statistical mechanics. In particular, general theoretical questions about strongly interacting quantum many-body states are starting to get answers and feedback from experimental efforts. On a more practical level –apart from this control leading to a better fundamental understanding of many-body physics– this may also lead to new developments in using quantum many-body physics as components of devices.

Although in this part of the thesis we are mainly concerned with the study of systems out-of-equilibrium, either by considering the dynamics of the state as it obtains equilibrium in the next chapter or by expliclity including reservoirs as in the previous chapter, we make a short exception to this theme and focus on closed quantum systems here. In a closed quantum system, the dynamics is fully governed by unitary evolution following the Schrödinger equation. Starting the system out in some general state that is not an eigenstate, we may ask the question: what is the nature of the long-time steady-state of the system?

For any finite-size system, the system will periodically return to its initial state (Poincaré recurrences). Practically, this timescale is exponential in the system size and hence is prohibitively long. In mathematically rigorous terms however, we should consider the thermodynamic limit to avoid such issues alltogether. In this limit, it is currently believed that there are two types of quantum many-body states: those that thermalize and those that do not [14, 17, 99–107]. The non-thermalizing states are also referred to as localizing states. Let us define in more detail what is meant by thermalizing and localizing.

5.1.1. THERMALIZING STATES

If we couple a system to an external reservoir at a given temperature, the system will eventually equilibrate with this reservoir. The system is then described by a thermal density matrix at the temperature of the reservoir. The observables in this state attain values that can be computed using a Gibbs ensemble corresponding to the reservoir temperature.

For the case of the *closed* system that we discuss here, there is no external reservoir and thermalization has to be understood as the system *itself* acting as a bath. In other words, a closed quantum system is called thermalizing if the dynamics of the system is such that a *subsystem sees the rest as an effictive thermal reservoir*. If this is indeed the case, statistical mechanics 'holds' and we can use any of the statistical ensembles to compute observables on the *full* system and get the correct results for the thermal subsystems.

However, there is an apparent paradox at hand. No matter whether we have equivalence of the statistical ensembles to compute expectation values, the closed system is still governed by unitary evolution. Unitary evolution cannot erase information about the initial

state, whereas a thermal state implies complete loss of the knowledge of the initial state. The resolution to this problem is found by realizing that the unitary evolution is able to spread, or 'wash out', this initial knowledge. In the long-time limit, this information becomes so much spread out that it would require a measurement of the full system at once (i.e. using a global operator) to retrieve it. If we were to consider only a subsystem, the initially local information is quickly lost to the entire system, which is effectively a large reservoir. From the point of view of the subsystem, the initial information has decohered.

Thermalization is a property that may in princple depend on the initial state of the system. It is not clear whether for a given system it is possible to have some initial states that would thermalize whereas some (possible special) initial states would not. But let us assume that we have a system for which indeed any initial state thermalizes. In particular, that means each of the eigenstates of the Hamiltonian thermalizes. And since eigenstates have trivial time-evolution that means these eigenstates themselves must be thermal. This is known as the strong version of the eigenstate thermalization hypothesis (ETH). A weaker version states that *almost* all eigenstates are thermal.

We can specify what it means for an eigenstate to be thermal in slightly more detail. Let us consider an eigenstate $|n\rangle$ of a many-body Hamiltonian H at energy E_n, i.e. $H|n\rangle = E_n|n\rangle$. Let us denote by $\langle H\rangle_T$ the expectation value of the Hamiltonian as measured in an actual thermal Gibbs state at temperature T. Then this eigenstate is a thermal state at a temperature defined through $E_n = \langle H\rangle_T$, and observables measured in this state should be reproduced by a Gibbs ensemble at this respective temperature.

The ETH is, as the name suggests, a hypothesis. It is numerically very challenging to check all eigenstates of an exponentially large many-body Hamiltonian, and hence it is not sure that the strong version of ETH holds in general. What is clear however, is that there is a class of systems for which the eigenstate thermalization hypothesis is definitely violated: the localized systems. Before we consider those systems, we mention that there is another way to avoid ETH. This is the case for integrable systems, which have an extensive set of conserved quantities. Unitary dynamics preserves these quantities, and so the system is constrained in its thermalization. It will thermalize all the degrees of freedom possible, leaving the final state in a generalized Gibbs ensemble. Integrability is usually a very special and fine-tuned point, and it is practically impossible to experimentally satisfy such exact conditions. The localizing states do not need such fine-tuning, and therefore seem to be the more robust candidates for avoiding eigenstate thermalization.

5.1.2. LOCALIZING STATES

Localizing states are those that by definition do not thermalize. The majority of studies that deal with localization are those about single-particle Anderson localization. Anderson introduced this concept for systems that have quenched disorder [16], meaning static disorder that is hard-coded in the system/Hamiltonian. We will briefly discuss the case of single-particle localization as a prerequisite for studying many-body localization.

Single particle Anderson localization is elegantly demonstrated by considering a tight-binding model of particles hopping on a disordered potential lattice:

$$H = \sum_{\langle ij \rangle} t\left(a_i^\dagger a_j + a_j^\dagger a_i\right) + \sum_i V_i a_i^\dagger a_i. \tag{5.1}$$

In this Hamiltonian, $a_i^\dagger$ creates a particle on lattice site i, $t \neq 0$ is the hopping amplitude between nearest neighbors (indicated using the notation $\langle ij \rangle$) and V_i is the potential strength on site i. We consider the case in which the potential values V_i are drawn from a random distribution characterized by a width, that we refer to as the disorder strength. In one or two spatial dimensions, for any non-zero disorder strength all of the eigenstates of this single particle Hamiltonian are exponentially localized with a localization length ξ that depends on disorder strength and energy [108]. This situation is demonstrated in Fig. 5.1, where we show a sample of the eigenstates $|\psi^n\rangle$ (i.e. $H|\psi^n\rangle = \epsilon_n|\psi^n\rangle$) of the Hamiltonian in Eq. 5.1 for a system with $L = 200$ sites. In three or more dimensions localization may still occur, but only for disorder strengths that are larger than some critical value [108]. For disorder strengths that are weaker, the states remain delocalized. Additionally, there is the option of having a mobility edge; an energy below which states are localized and above which they are delocalized. The states that are exactly at the mobility edge are also localized, but decay much slower following a power-law instead of exponential decay. To construct the many-body eigenstates of the Hamiltonian in

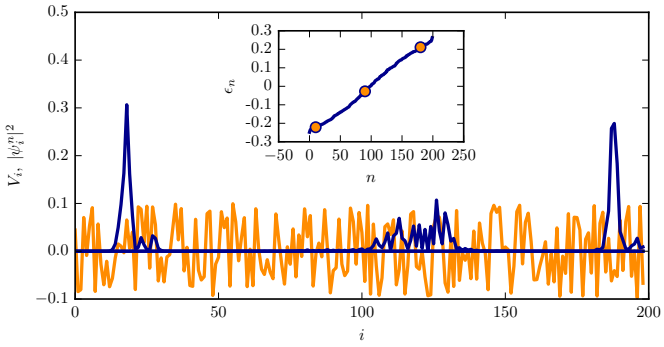

Figure 5.1: Anderson localization demonstrated on the non-interacting disordered system of Eq. 5.1 on 200 sites. Shown are three eigenstates at different energies (in blue), all of them localized in the disorder potential V_i (in orange).

Eq. 5.1 we may just consider product states of the single particle eigenstates $|\alpha\rangle$, assigning to them occupation numbers n_α. If we consider a system in which some of the single particle eigenstates are localized, it turns out many of the many-body eigenstates violate the eigenstate thermalization hypothesis [17]. The typical example of this is given by an initial state in which we construct a non-uniform density distribution. The single particle localized eigenstates at the respective spatial positions prevent the density pattern from becoming uniform and make for a persistive density pattern.

We remark that localization as we have just discussed is not restricted to particles hopping on a lattice. In fact, we will be considering mostly a spin system. In the single particle sector, a direct comparison is straightforward: a single particle on an empty lattice is equivalent to a single spin-↑ on a background of only spin-↓'s.

Let us turn now to localization in the presence of interactions, i.e. true many-body localization (MBL). In a many-body localized system, the localization happens at all energy scales. In particular, we will be discussing high-energies to exclude single particle localization effects at lower energies, and to distinguish it from general 'glassy' behaviour. Additionally, this means that we don't have to worry about distinguishing bosons from fermions from spins, and we will be working mostly with the latter in accordance with most of the literature [17]. The typical model we study is given by the Hamiltonian:

$$H = \sum_i J\boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_{i+1} + \sum_i h_i \sigma_i^z, \tag{5.2}$$

where the Pauli operators $\boldsymbol{\sigma}$ describe spin-1/2 degrees of freedom with nearest neighbor coupling J, and an external magnetic field h_i. The external field h_i is here the component that provides the quenched disorder.

Arguments from perturbation theory around small interactions J shows that at least for weak enough interactions the eigenstates still violate ETH. Or in other words, if we fix $J = 1$ we may say that for strong enough disorder the eigenstates violate ETH. There is considerable numerical evidence that this is indeed the case, and remains true even interaction strengths that are stronger than perturbation theory allows them to be. For very strong interactions, eigenstates start obeying ETH and the system delocalizes.

The model in Eq. 5.2 shows a quantum phase transition between localizing and thermalizing behaviour as a function of disorder strength. This transition happens in the eigenstates, and is 'invisible' to equilibrium thermodynamic observables. Hence we have to necessarily study out-of-equilibrium dynamics to learn about the transition. Examples of this are coupling the MBL system to external baths, or by studying transient behaviour after starting the system in an initial state that is not an eigenstate (i.e. equivalent to a quench). In the following section we will discuss the decoherence properties of a single probe spin coupled to an environment described by Eq. 5.2, and find that indeed the different regimes lead to different properties.

Before turning to the decoherence however, we review a phenomenological approach to MBL that will turn out to provide a useful interpretation of the results. The nomenclature that exists in the literature revolves around p-bits, for physical-bits, by which we denote the actual physical degrees of freedom. In the Hamiltonian of Eq. 5.2 this would be the $\boldsymbol{\sigma}$ operators. Deep in the MBL phase, one may argue [19, 106, 109–112] that the Hamiltonian can instead be written as

$$H = E_0 + \sum_i \tau_i^z + \sum_{ij} J_{ij}\tau_i^z\tau_j^z + \sum_{n=1}^{\infty}\sum_{ij\mathbf{k}} K_{i\mathbf{k}j}^{(n)}\tau_i^z\tau_{k_1}^z\ldots\tau_{k_n}^z\tau_j^z. \tag{5.3}$$

In this Hamiltonian, we have introduced a new set of Pauli operators $\boldsymbol{\tau}$ that we refer to as l-bits, for localized-bits. The sums in Eq. 5.3 are such that each term appears only once,

and E_0 is an irrelevant energy offset. The essential point of this transformation is that deep within the MBL phase the magnitudes of the interactions J_{ij} and $K^{(n)}_{i\mathbf{k}j}$ fall of exponentially with distance. Notice that this Hamiltonian only consists of the τ^z operator, and hence commutes with τ^z on every site. Hence we have an extensive set of conserved operators as is the case for an integrable system.

Intuitively, the presence of such localized-bits is due to the absence of conduction in a localized phase. In the context of long-time conserved quantitities, these l-bits are also referred to as local integrals of motion (LIOMs). We may deduce an important property of the MBL phase from the construction of these l-bits, namely the observed logarithmic growth of entanglement entropy. In a thermalizing system entanglement spreading is usually ballistic [113, 114], whereas in a single-particle localized system there is no entanglement spreading at all. The single particle localized case is namely the extreme case of Eq. 5.3 in which the interactions terms (J and K) between the l-bits are absent. In an MBL system, where J and K are present but decay exponentially, the l-bits can entangle via direct interactions. This latter point is again in contrast to thermalizing systems in which a region A may become entangled with region C, if both A and C talk to a region B. For the MBL case, such situations do not occur since the interactions are only diagonal, i.e. they do not involve spin flips. So an l-bit A cannot change the l-bit B, and an l-bit C that also talks to B would therefore not be influenced by what A does to B.

This can be quantified a bit better, by considering the effective interaction between two l-bits i and j as

$$J^{\text{eff}}_{ij} = J_{ij} + \sum_{n=1}^{\infty} \sum_{\mathbf{k}} K^{(n)}_{i\mathbf{k}j} \tau^z_{k_1} \dots \tau^z_{k_n}, \tag{5.4}$$

so that the interaction depends on all the l-bits in between i and j. Due to the exponential decay of J and K, we argue that $J^{\text{eff}} \sim J_0 e^{-L/\xi}$. If we consider only these two l-bits that are initially unentangled, it takes a time of order $t \leq 1/J^{\text{eff}}$ for them to get entangled. So in a time t, we can entangle l-bits that are at a distance $L \sim \xi \ln J_0 t$ from each other. This logarithmic dependence is an important signature of the MBL phase [115].

5.2. SINGLE PROBE DEPHASING

Parts of this section have been published in **E.P.L. van Nieuwenburg, S.D. Huber and R. Chitra**, *Single spin probe of Many-Body Localization*, arXiv e-prints 1607.08617.

5.2.1. MOTIVATION

The state of an object in a classical world can be represented by e.g. a thermal state, as is the case in a statistical mechanics description. In the language of density matrices, this means that we only assign probabilities to the system being in some given configuration (with the probability depending on the energy of the configuration through the

Boltzmann factor). In contrast to such diagonal density matrices, a quantum system is characterised by having non-zero off-diagonal entries, also referred to as *coherences*. The loss of coherence then, is what signals the transition from the quantum world to the classical world. Exactly how a quantum system decoheres is therefore a question of significant importance, both for our understanding of the fundamental princples as well as for potential applications that try to make use of this coherence.

In this section we will turn the question around, asking what we can learn about the system given its decoherence properties. We will apply this reasoning to distinguish the many-body localized phase from the thermalizing phase. The idea of studying the decoherence properties has for example been applied to finding out whether a system is a thermal bath or not [116–118].

Instead of considering a system with an MBL transition and coupling it to a bath to study its decoherence, we use the MBL system as a bath itself. The system will act as a bath for a single spin that we couple to it, and whose decoherence we will be concerned with. This is fundamentally different from studying the MBL system itself, since while the failure of the ETH is a fascinating phenomenon, it entirely relies on the perfect isolation of the system. Any globally accessible heat-bath will necessarily equilibrate the system and render the notion of ETH obsolete.

Because in experiments a heatbath is always present to some degree, the experimental realisation of an MBL system represents a difficult challenge [119–121]. Regardless, recent experiments based on cold atoms [104, 122, 123] and trapped ions[105] have tackled this challenge, and showed that it is possible to retain the order of the initial state on long timescales. Moreover, it is clear from the experiments that it is the coupling to the environment that limits this timescale. We will next develop a minimal probe for an MBL system based on the single spin that we couple it to.

Our setup is sketched in Fig. 5.2, with the circles representing spin-1/2 degrees of freedom. The last spin is considered to be an external probe spin, and is coupled to the system in a specific way (to be described later). Instead of observing or measuring the system, we will deduce its properties solely by studying the decoherence dynamics of the single spin attached to it. This idea in general has a rich history [124–132]. Especially important to this chapter, is the result that the level statistics of the bath influence the decoherence properties of the single spin [133]. We will see that level statistics is able to capture the MBL transition, and we therefore expect to see the effect of this in the decoherence of the probe.

5.2.2. MODEL

The MBL reservoir that we study is a Heisenberg spin chain in a random field, already presented in Eq. 5.2, but repeated here for convenience:

$$H_s = J \sum_{i=1}^{L-1} \boldsymbol{S}_i \cdot \boldsymbol{S}_{i+1} + \sum_{i=1}^{L} h_i S_i^z, \tag{5.5}$$

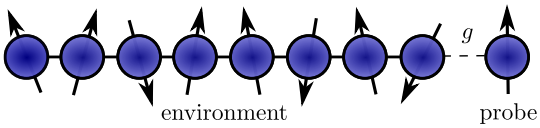

Figure 5.2: Sketch of the setup. A single probe spin is coupled to (the end of) a disordered and interacting spin-chain forming the environment. The probe is coupled in such a way that its dynamics is governed purely by dephasing, which is strongly influenced by the state of the environment.

where $J > 0$ is the exchange coupling, $\boldsymbol{S}$ are the canonical spin operators and the random fields h_i are independently and uniformly distributed in the interval $[-h_{\max}, h_{\max}]$. Let us list the key properties of this system in a bit more detail than we did in section 5.1.2. Most importantly, this model is known to exhibit a transition between a many-body localized phase and a thermalizing phase as a function of $h_{\max}$. Previous studies [75, 100] indicate that in the thermodynamic limit the MBL transition occurs at $h_{\max} \sim 3.5$. This means that the closed system obeys ETH and thermalizes if the disorder strength $h_{\max}$ is lower thatn 3.5. If disorder strength is larger than this, the system is many-body localized, violates ETH and hence avoids thermalization.

The absence of thermalization in this model is not due to fine-tuning to integrability but arises from an emergent integrability. This integrability is encoded in the l-bits of section 5.1.2, but we will referred to them rather as (quasi-) local integrals of motion (LIOM's) [19, 106, 109, 110].

The transition between the thermalizing and non-thermalizing regime is visible in the statistics of the eigenvalue distribution (also called level spacing statistics) [18, 134, 135]. More precisely, the average level-spacing between eigenvalues evolves from a Wigner-Dyson distribution at low values of $h_{\max}$ (in the thermalizing phase), to a Poissonian distribution in the MBL phase. In the next subsections, we will illustrate how these results come about. In particular, for the Wigner-Dyson distribution we turn to random matrix theory.

WIGNER-DYSON DISTRIBUTION

As an example, let us consider a Hamiltonian H for a disordered system that has only real entries. For simplicity, we consider a two-level system so that such a Hamiltonian is a random 2×2 matrix. It then belongs to the so called gaussian orthogonal ensemble (GOE) random matrices, characterized by H being real and satisfying $H = H^T$. We will denote the spacing between its two eigenvalues as Δ, and prove below that the probability distribution $P(s)$ for Δ is given by

$$P(s) = \frac{\pi s}{2} e^{-\frac{\pi s^2}{4}}. \tag{5.6}$$

To do this , we first define a probability distribution on the level of the Hamiltonian (i.e. for random matrices) as

$$P(H) = \mathcal{N} e^{-\mathrm{Tr}H^2}, \tag{5.7}$$

where $\mathcal{N}$ is a normalization constant. More precisely, this equation defines the joint probability distribution function of all the independent entries in H. Using this, we may define the following notation for computing an expectation value:

$$\mathbb{E}(f) = \int \mathrm{d}H f(H) P(H). \tag{5.8}$$

For the distribution function of Eq. 5.6, we have assumed a normalization such that $\mathbb{E}(\Delta) = 1$. This is a condition that can always be realised on a given set of eigenvalues, but in general (i.e. for more than just 2×2 matrices) requires the 'unfolding' of the spectrum [136]. This is not a difficult procedure, but we will assume for simplicity that this has already been achieved.

Since H is a matrix that belongs to the GOE class, we can parameterize it as

$$H = \begin{pmatrix} a & c \\ c & b \end{pmatrix},$$

with a, b, c being real. The integration measure $\mathrm{d}H$ is then given simply by $\mathrm{d}a\,\mathrm{d}b\,\mathrm{d}c$.

The two eigenvalues $\lambda_\pm$ of H are

$$\lambda_\pm = \frac{1}{2}\left[(a+b) \pm \sqrt{(a-b)^2 + 4c^2}\right],$$

meaning that

$$\Delta = \sqrt{(a-b)^2 + 4c^2}.$$

The probability function $P(s)$ is now straightforwardly calculated as

$$\begin{aligned} P(s) &= \mathbb{E}(\delta(s-\Delta)) \\ &= \mathcal{N} \int \mathrm{d}H \delta(s-\Delta) e^{-\mathrm{Tr}H^2} \\ &= \mathcal{N} \int \mathrm{d}a\,\mathrm{d}b\,\mathrm{d}c \delta(s - \sqrt{(a-b)^2 + 4c^2}) e^{-a^2-b^2-2c^2} \\ &= \mathcal{N} \int_{-\infty}^{\infty} \mathrm{d}a e^{-a^2} \int_{-\infty}^{\infty} \mathrm{d}b e^{-b^2} \int_{-\infty}^{\infty} \mathrm{d}c e^{-c^2} \delta(s - \sqrt{(a-b)^2 + 4c^2}). \end{aligned} \tag{5.9}$$

This integral may be solved by substituting $x = a - b$ and $y = a = b$ and switching to polar coordinates. The result is

$$P(s) = s e^{-s^2/2},$$

where we have taking into account the normalization factor $\mathcal{N}$. Using this result, we may calculate $\mathbb{E}(\Delta)$ and find out that it equals $\sqrt{\pi}2$ and is not equal to unity. In the present case of 2×2 matrices, making $\mathbb{E}(\Delta) = 1$ can be done by a simple rescaling:

$$\tilde{P}(s) = \sqrt{\pi}2 P(s\sqrt{\pi}2) = \frac{\pi s}{2} e^{-\frac{\pi s^2}{4}}.$$

This result for the distribution of the energy level spacing can be extended to random Hamiltonians with complex entries, belonging to the gaussian unitary ensemble (GUE). It is refered to as the 'Wigner-surmise', or also as the Wigner-Dyson distribution.

POISSONIAN DISTRIBUTION

Let us instead think of a Hamiltonian whose eigenvalues are completely uncorrelated. Physically, such a situation may indeed arise when its eigenstates are perfectly localized and hence do not hybridize. In this case, the probability that an eigenvalue ϵ_n will lie within the interval ϵ and $\epsilon + \mathrm{d}\epsilon$ is indepdendent of ϵ, and is given by $D^{-1}\mathrm{d}\epsilon$ with D the number of energy values within a unit of energy (i.e. the density of states). The probability distribution $P(s)$ in this case is now interpreted as: given an eigenvalue at energy ϵ, what is the probability of having no eigenvalue in the interval $(\epsilon, \epsilon + s)$ and a single eigenvalue in the interval $(\epsilon + s, \epsilon + s + \mathrm{d}s)$?

The trick to answering this question, is by dividing the interval between ϵ and $\epsilon + s$ into m equal parts. Due to the levels being uncorrelated, the probability of having no level in $(\epsilon, \epsilon + s)$ is the product of having no level in any of the smaller m intervals. For a given interval with small m this probability is given by $1 - D^{-1}\frac{s}{m}$, and hence for the full interval with $m \to \infty$:

$$\lim_{m\to\infty}\left(1 - D^{-1}\frac{s}{m}\right)^m = e^{-D^{-1}s}.$$

Hence the probability of no level in $(\epsilon, \epsilon + s)$ and one level in withing $\mathrm{d}s$ at $\epsilon + s$ is the product $e^{D^{-1}s}D^{-1}\mathrm{d}s$. We define a rescaled average level spacing $s \to D^{-1}s$, and find that the average level spacing is hence distributed as

$$P(s)\mathrm{d}s = e^{-s}\mathrm{d}s. \tag{5.10}$$

This is the Poisson distribution for the level spacing of uncorrelated energy levels.

R-STATISTIC

Instead of calculating the full probability distributions $P(s)$, it is worthwhile to characterise both distributions by a single number. There are many characteristics, but a convenient [18] characteristic we will consider is the r-statistic, since it does not suffer from having to unfold the spectrum.

We start by defining the ratio of consecutive level spacings

$$r^{(n)} = \frac{\min\left(s^{(n)}, s^{(n+1)}\right)}{\max\left(s^{(n)}, s^{(n+1)}\right)},$$

where $s^{(n)} = \epsilon_n - \epsilon_{n+1}$. We may now average this quantity over the full spectrum, or as is commonly done only over the bulk of the spectrum to prevent edge effects, and denote it r. The r-statistic is then defined simply as the disorder average of r, denoted $\langle r \rangle$. For

a level spacing distribution belonging to the GOE ensemble, $\langle r\rangle \approx 0.53$ [137], whereas for the Poissonian distribution it may be calculated exactly as $\langle r\rangle = 2\ln 2 - 1 \approx 0.39$. The r-statistic for our model is given in panel **A** of Fig. 5.3.

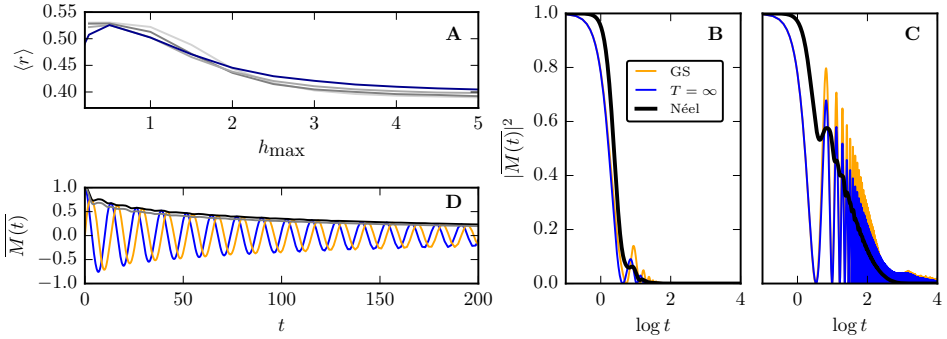

Figure 5.3: **A**: The level spacing statistics for $L = 8$ (blue) up to $L = 14$ (lightest grey) in steps of two, summarized by the $\langle r\rangle$ statistic [18]. For a Wigner-Dyson distribution $\langle r\rangle$ approaches the value ~ 0.53, whereas for the Poissonian distribution $\langle r\rangle$ approaches ~ 0.39. The estimate $h_{\max} \approx 3.5$ for the transition point can be obtained by interpolating the crossing points of $\langle r\rangle$ for successively larger system sizes. **B and C**: $|\overline{M(t)}|^2$ for three initial states of the environment: the groundstate (GS, orange), infinite temperature state ($T = \infty$, blue) and the Néel state (black). The left panel shows results for $h_{\max} = 1.0$ whereas the right panel shows results for $h_{\max} = 5.0$. From the $\langle r\rangle$ statistic, we know that these are two extremes in the thermalizing and MBL regime, respectively. In the latter, the coherence is retained on much longer timescales for all of the initial states. The Néel state is a convenient choice of state, as it basically extracts the shape of the envelope of the decay. **D**: The real (blue) and imaginary (orange) parts of $\overline{M(t)}$ for an initially Néel ordered environment with $h_{\max} = 5$. The black line shows $|\overline{M(t)}|$ while the gray bar shows $|M(t)|$ as obtained from the entropy (see main text).

Let us get back to the description of the decoherence properties of the probe spin. Formally, the Hamiltonian of the full system reads

$$H = H_{\rm s} + H_{\rm ep}, \tag{5.11}$$

where $H_{\rm s}$ is given in Eqn. (5.5) and $H_{\rm ep}$ describes the environment-probe interaction, given by

$$H_{\rm ep} = 2\sum_i g_i S_i^z \sigma^z. \tag{5.12}$$

Here, the σ^a ($a \in \{x, y, z\}$) operators describe the probe spin. Even though this spin is not physically different from those in the system, we choose a different label for clarity. Notice that there is no term Hamiltonian describing the probe spin by itself. Such a term would not influence our discussion below, and since we are fundamentally interested in the decoherence rather than any relaxation, we have left it out.

The probe spin is coupled to the end of the system, and hence we set $g_L = g \neq 0$ and all other g_i are to zero. As opposed to coupling the spin to a large fraction of the system,

which is not particularly sensitive to localization [132], by coupling the test-spin to a single site we will be susceptible to localization and the associated emergent integrability of the MBL phase. Let us now discuss this decoherence dynamics.

5.2.3. DYNAMICS

The probe spin is initialized in a pure state $\rho = |\psi_0\rangle\langle\psi_0|$ with $|\psi_0\rangle = \alpha|\uparrow\rangle + \beta|\downarrow\rangle$. We will defer the exact choice of the state of the environment, but mention that it is described by a density matrix ρ_e. The combined system is then assumed to be in a factorizable initial state at time $t = 0$, denoted $\Omega = \rho_e \otimes \rho$. The combination of this initial state with the type of coupling in Eq. 5.12 leads to a pure decoherence of the probe spin. This is because the coupling term does not induce spin-flips, and hence leaves the diagonal entries of ρ untouched.

At time $t = 0^+$ we turn on the coupling, and let the system undergo the usual unitary time evolution governed by the Hamiltonian. Let us consider the off-diagonal element of the probe density matrix:

$$\begin{aligned}\langle\downarrow|\rho(t)|\uparrow\rangle &= \langle\downarrow|\mathrm{Tr}\Omega(t)|\uparrow\rangle \\ &= \mathrm{Tr}\left(\langle\downarrow|e^{-i(H+gS_L^z\sigma^z)t}(\rho_e\otimes\rho)e^{i(H+gS_L^z\sigma^z)t}|\uparrow\rangle\right) \\ &= \mathrm{Tr}\left(e^{-i(H-gS_L^z)t}\rho_e e^{i(H+gS_L^z)t}\right)\langle\downarrow|\rho|\uparrow\rangle \\ &= M(t)\alpha^*\beta \end{aligned} \tag{5.13}$$

where the trace is understood to be taken over the environment, and we have defined the factor $M(t)$ as

$$M(t) = \mathrm{Tr}\left(e^{-i(H-gS_L^z)t}\rho_e e^{i(H+gS_L^z)t}\right). \tag{5.14}$$

This factor is the measure for the coherence of the probe spin at time t, since α and β remain unchanged. For completeness, we include the full time dependent density matrix of the probe:

$$\rho(t) = \begin{pmatrix} |\alpha|^2 & M(t)\alpha^*\beta \\ M(t)^*\alpha\beta^* & |\beta|^2 \end{pmatrix}. \tag{5.15}$$

In the weak-coupling limit, for which g is much smaller than both the disorder strength $h_{\max}$ and J, the Born approximation that the bath is uninfluenced by the probe holds. In this limit, the coherence factor $M(t)$ is essentially determined by the end spin correlator $\mathrm{Tr}\left(\rho_e\sigma_L^z(t)\sigma_L^z(0)\right)$ [126]. We will not assume that the Born approximation holds, instead fixing the coupling strength to $g = J = 1$. Then, we resort to the use of numerical exact diagonalization to obtain $M(t)$ for many different disorder realizations. Additionally, we mention that we restrict ourselves to the total $\sum_{i=1}^{L}\sigma_i^z = 0$ sector. This restriction allows us to directly study the level statistics of the model, and is a physically relevant constraint in experiments. The fixed spin-sector corresponds to a fixed particle number, as is the case for example in cold atomic systems simulating spins.

We will be mostly concerned with the disorder average of $M(t)$, which we denote with an overbar as $\overline{M(t)}$. Whereas $M(t)$ never decays for single realisations, due to the finite size recurrences, as time $t \to \infty$ the coherence $\overline{M(t)}$ *does* go to zero. We therefore focus on the intermediate time scales where the coherence is not yet lost, and we will see that this transient regime of $\overline{M(t)}$ depends on the properties of the environment. In particular, this transient regime is sensitive to whether the system is thermalizing or many-body localized.

Finally, before we turn to the results, let us shortly discuss the choice of initial state for the environment. We consider three different initial states:

- The groundstate for each given disorder realization
- The infinite temperature state, described by a density matrix $\rho_e = \mathbb{1}/2^L$
- A high-energy Néel ordered state, in analogy with experiments [104].

Especially the infinite temperature state and the excited Néel ordered state are interesting initial states, since they emphasize the MBL aspect of violating ETH also in high-energy states. For reasons that we discuss below, we will restrict ourselves to the discussion of the Néel state.

5.2.4. RESULTS

The use of exact diagonalization limits the system sizes we can simulate with the available computing power (and within reasonable timescales) to about 12-14 sites, since we require the full spectrum of the Hamiltonian. Empirically, we find that already for a system of size $L = 8$ and about 10^6 disorder realizatoins we observe the main features of the r-statistic (see panel **A** of Fig. 5.3). As we will discuss later, larger system sizes hardly have any influence on our results in the first place.

Panels **B** and **C** of Fig. 5.3 show results for the disorder averaged coherence of the three choices of initial states, at disorder strengths deep in the delocalized ($h_{\max} = 1$) and localized ($h_{\max} = 5$) phases. Each of the choices of initial state shows markedly different behaviour between the small and large disorder strength results. Most notably, the coherence is lost on a much shorter timescale in the delocalized regime than it is in the MBL regime.

Another feature that is visible is the pronounced ringing in the MBL regime. This can be understood by considering the LIOMs. In the MBL regime, the probe spin couples mostly to the LIOM at the end of the bath, and hence performs coherent Rabi-oscillations at a frequency set by the coupling g. This frequency is dressed slightly due to the presence of other spins in the environment, but is mostly constant and *does not depend on the particular realisation of disorder*. Therefore, the averaging will not wash out this ringing behaviour.

Also clearly visible in panel **C** is the suppression of the dips in the ringing for the state that has Néel order. It essentially captures the envelope of the $T = \infty$ state, up to a rescaling of the time axis. Therefore, in the following we focus exclusively on the Néel state and analyze the different time regimes.

Without regards to the timescales yet, we remark that the real and imaginary parts of $\overline{M(t)}$ can be measured separately, and correspond to the expectation values of the probe spin in the x- and y-direction respectively. An example for the Néel state at large disorder strength is shown in panel **D** of Fig. 5.3. The resulting absolute value $|\overline{M(t)}|$ is shown in black. Additionally, the gray curve shows the coherence factor obtained from the entropy of the probe. This is to show that really the information in $M(t)$ also allows for statements on the entropy, which we know is a useful tool in many-body physics. To make the connection between $M(t)$ and the entropy $S(t)$, note that for small $|\overline{M(t)}|$ (i.e. in the long-time limit) the entropy $S(t) = -\mathrm{Tr}\rho(t)\ln\rho(t)$ of the probe spin can be calculated from an expansion as follows. The density matrix $\rho(t)$ for the probe spin can be written as $\rho(t) = \frac{1}{2}\left(\hat{\mathbf{1}} + \boldsymbol{d}\cdot\boldsymbol{\sigma}\right)$, where the vector $\boldsymbol{d}$ is given by

$$\boldsymbol{d} = \left\{2\,\mathrm{Re}\,\alpha\beta^* M(t), 2\,\mathrm{Im}\,\alpha\beta^* M(t), |\alpha|^2 - |\beta|^2\right\}.$$

Denoting by $|d|$ the norm of this vector, we have in particular that due to the Pauli matrices $(\boldsymbol{d}\cdot\boldsymbol{\sigma})^2 = |d|^2\hat{\mathbf{1}}$. The norm is straightforwardly calculated to be

$$|d|^2 = 1 - 4|\alpha|^2|\beta|^2\left(1 - |M(t)|^2\right). \tag{5.16}$$

We will drop the explicit time-dependence in the following, for the sake of clarity. The logarithm of ρ can be expanded as

$$\begin{aligned}\ln\rho &= \ln\left(\frac{1}{2}\left(\hat{\mathbf{1}} + \boldsymbol{d}\cdot\boldsymbol{\sigma}\right)\right)\\ &= -\ln 2\,\hat{\mathbf{1}} + \sum_{n=1}\frac{(-1)^{n+1}}{n}(\boldsymbol{d}\cdot\boldsymbol{\sigma})^n,\end{aligned} \tag{5.17}$$

which means that, using the tracelessness of the Pauli matrices, the entropy can be written as

$$\begin{aligned}S &= -\mathrm{Tr}\,\rho\ln\rho\\ &= -\frac{1}{2}\mathrm{Tr}\ln\rho - \frac{1}{2}\mathrm{Tr}\left(\boldsymbol{d}\cdot\boldsymbol{\sigma}\ln\rho\right)\\ &= \ln 2 - \frac{1}{2}\sum_{n=1}\frac{(-1)^{n+1}}{n}\mathrm{Tr}\left((\boldsymbol{d}\cdot\boldsymbol{\sigma})^n\right) - \frac{1}{2}\sum_{n=1}\frac{(-1)^{n+1}}{n}\mathrm{Tr}\left((\boldsymbol{d}\cdot\boldsymbol{\sigma})^{n+1}\right).\end{aligned} \tag{5.18}$$

Again due to the tracelessness of the Pauli matrices the terms with n odd do not contribute to the first sum, while those with n even do not contribute in the second sum. Together with the fact that $\mathrm{Tr}\left((\boldsymbol{d}\cdot\boldsymbol{\sigma})^n\right) = 2|d|^n$ for even n, we find:

$$\begin{aligned}S &= \ln 2 - \sum_{n=1}\frac{(-1)^{2n+1}}{2n}\left(|d|^2\right)^n - \sum_{n=0}\frac{(-1)^{2n+2}}{2n+1}\left(|d|^2\right)^{n+1}\\ &= \ln 2 - \frac{1}{2}\ln\left(1 - |d|^2\right) - |d|\tanh^{-1}|d|.\end{aligned} \tag{5.19}$$

In particular, at $t = 0$ we know that $M(t) = 1$ and hence $|d|^2 = 1$ and we correctly find that $\lim_{d\to 1} S = 0$ for the pure state. Let us specify to the case $\alpha = \beta = 1/\sqrt{2}$, which maximizes the product $|\alpha|^2|\beta|^2$ under the constraint that $|\alpha|^2 + |\beta|^2 = 1$. In this case $|d| = |M(t)|$, and hence for small $|M(t)|$ we may expand the result for S as:

$$\begin{aligned} S &= \ln 2 + \frac{1}{2}|M(t)|^2 - |M(t)|^2 \\ &= \ln 2 - \frac{1}{2}|M(t)|^2 \end{aligned} \tag{5.20}$$

So in the long-time limit, $S \to \ln 2$ as expected for a singlet. In general, this limit will depend on the choice for α and β. The gray curve in Fig. 5.3 shows $|\overline{M(t)}|$ obtained from the entropy in this way. For times $t > 10$, the two curves are essentially identical.

Let us return to discussing the decoherence for the initial Néel ordered states. On short timescales ($0 < t \lesssim J^{-1}$) we may expand Eq. 5.14 to second order in t. For notational clarity we write $V = gS_L^z$ for the interaction term, and obtain

$$\begin{aligned} M(t) &= \mathrm{Tr}\left[e^{-i(H_\mathrm{s}+V)t}\,\rho_\mathrm{e}\,e^{i(H_\mathrm{s}-V)t}\right] \\ &\approx \mathrm{Tr}\left[\left(1 - i(H+V)t - \frac{1}{2}(H+V)^2t^2\right)\left(1 - i(H-V)t - \frac{1}{2}(H-V)^2t^2\right)\right] \\ &= 1 - 2it\langle V\rangle - 2t^2\left(\langle V^2\rangle - \frac{1}{2}\langle [H,V]\rangle\right) \end{aligned} \tag{5.21}$$

Hence in $|M(t)|^2$ the lowest order term is of the form $2t^2\langle [H,V]\rangle$, leading to Gaussian decay $|M(t)|^2 \sim e^{-(t/t^*)^2}$ for all disorder strengths, with t^* having only a very weak dependence on h_max. This Gaussian behaviour defines the initial time scale, and is a universal behaviour. After this initial decay period, deviations that depend on the disorder strength start to arise. Especially when comparing the weak ($h_\mathrm{max} \lesssim 3.5$) and strong ($h_\mathrm{max} \gtrsim 3.5$) disorder regimes, it is clearly observable that the timescale for loss of coherence is much longer in the latter one (see panels **B** and **C** of Fig. 5.3). The longer timescale in the strong disorder regime can be understood by considering the properties of the environment in more detail. In this regime, it is possible to introduce (at least in principle) an extensive set of l-bits or local integrals of motion (LIOMs) [19, 106, 109–112]. The probe spin couples predominantly to the LIOM at the end of the environment, and since the overlap between LIOMs is exponentially small this leads to information diffusion through the system only on an exponentially large timescale. In turn, that means that the probe spin is mainly exchanging coherence with this LIOM at the end (i.e. performing Rabi-oscillations), ensuring that even after averaging the coherence is retained.

On the intermediate timescale $J^{-1} \ll t \ll t_f(h_\mathrm{max})$ during which most of the decoherence occurs, it is possible to rescale the time axis for the large disorder strength and make $|\overline{M(t)}|$ collapse onto one curve. The top panel of Fig. 5.4 shows this by rescaling the time axis $t \to t/h_\mathrm{max}^\alpha$ with an exponent $\alpha = 2$ that does not depend on the disorder strength. For lower values of disorder strength, such a rescaling is not possible.

Since we know that from the intermediate time onwards $S(t) \approx \ln 2 - \frac{1}{2}|M(t)|^2$, and that

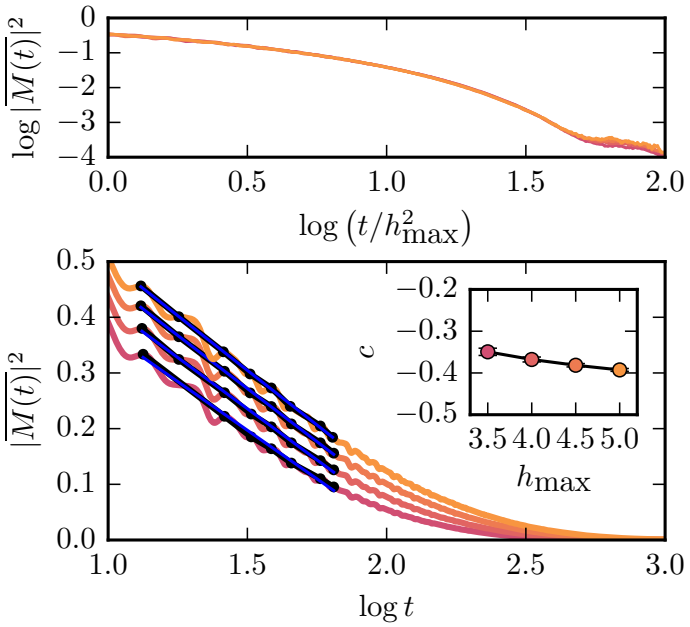

Figure 5.4: **Top panel**: In the MBL regime for $h_{\max} \gtrsim 3$ the curves for $\log|\overline{M(t)}|^2$ can be collapsed by rescaling $t \to t/h_{\max}^2$. This scaling works well in the intermediate timescale. **Bottom panel**: The intermediate times can be fit well by a logarithmic decay, $|\overline{M(t)}|^2 \sim c\log t$, with $c \approx -0.37$. This is in agreement with the observed logarithmic growth of the entanglement entropy in the MBL phase.

for the MBL regime $S(t) \sim c\log t$ [115], we propose to fit the intermediate timescales to a logarithmic decrease. Since $|M(t)|^2$ is bounded and must approach zero upon $t \to \infty$, this logarithmic growth is not expected to hold all the way to long timescales. As we show in the bottom panel of Fig. 5.4, we can indeed fit the coherence $|\overline{M(t)}|^2$ to $c\log(t)$ with c relatively constant constant close to -0.37 (shown in the inset of the panel). This observation is experimentally relevant, since it shows that the (relatively easy to measure) decoherence of the probe spin provides a direct way of measuring the logarithmic entanglement growth. Combined with the fact that from the rescaling of the curves we conclude $t_f \propto h_{\max}^2$, the logarithmic regime lasts for longer times as the disorder strength increases.

Another manifestation of the MBL regime can be found by considering the coherence of the probe spin at fixed moments in time, but as a function of disorder strength. We consider times in the intermediate regime, denoted τ, and plot $|\overline{M(\tau)}|$ as a function of $h_{\max}$ in Fig. 5.5. As can be seen in the figure, the scaling with disorder strength in the large disorder regime behaves differently from the small disorder regime. In the former, the slopes of $|\overline{M(\tau)}|$ with $h_{\max}$ are disorder strength independent, whereas the slops in the delocalized regime vary strongly. We have not investigated the asymptotic times $t \gg$

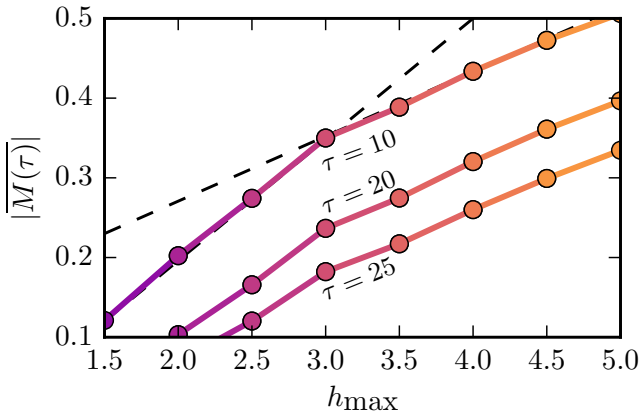

Figure 5.5: The values of $\overline{|M(\tau)|}$ at fixed times $\tau = 10, 20$ and 25 in the intermediate time regime, as a function of disorder strength $h_{\max}$. The dependence of $\overline{|M(\tau)|}$ on $h_{\max}$ shows a change of slope for the localized regime. The black lines are guides for the eye.

$t_f(h_{\max})$ in great detail. A new regime sets in where $\ln|M|(t) \propto -t^{\nu}$ where $\nu \approx 4$ and has a slight dependence on $h_{\max}$ for the MBL regime.

5.2.5. DISCUSSION

We have checked the influence of the coupling strength g on the results, presented in Fig. 5.6. In the MBL regime, the different curves for $\overline{|M(t)|}$ are related to eachother by a simple scaling of the form $t \to g\,t$. That means that all the features that we have discussed in this regime can be observed for both weak coupling as well as strong coupling. It also emphasizes that the MBL regime is localized and the coupling strength only determines the timescale over which the state is probed. For the smaller disorder strengths, a scaling by a factor g of time does not collapse the curves. However, since the decoherence in this regime is extremely fast, it is hard to observe and deduce the specific nature of this decay.

Additionally, we have checked our results against coupling of the spin in the middle of the chain. The result is that the timescales in the MBL regime shift by a factor of two, which is expected due to the probe effectively coupling to two systems. However, if we couple the probe at the end but increase system size, there is little to no effect (see Fig. 5.7). This, too, can be understood from the fact that the probe couples to the LIOM at the end. In the initial and intermediate timescales, the spin is hence unable to probe the size of the system.

To conclude, we have shown in this chapter that is is possible to identify whether a system is many-body localized or not, by observing the decoherence properties of a single spin probe coupled to it. This is possible by considering the scaling behaviour with re-

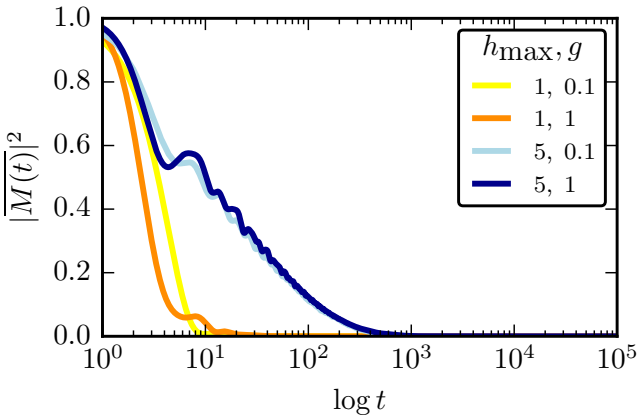

Figure 5.6: In the MBL regime, rescaling $t \to gt$ for different values of the coupling strength (shown here are $g = 1$ and $g = 0.1$, on a rescaled time axis) leads to a collapse of the curves for $|\overline{M(t)}|^2$. The low disorder regime does not obey such a scaling behaviour.

5

spect to disorder strength $h_{\max}$ as well as with respect to the coupling strength g. Moreover, the intermediate timescales show the characteristic logarithmic dependence on time expected for a many-body localized system.

Our setup is possibly realizable in ion trap systems, where the coherence of a probe spin

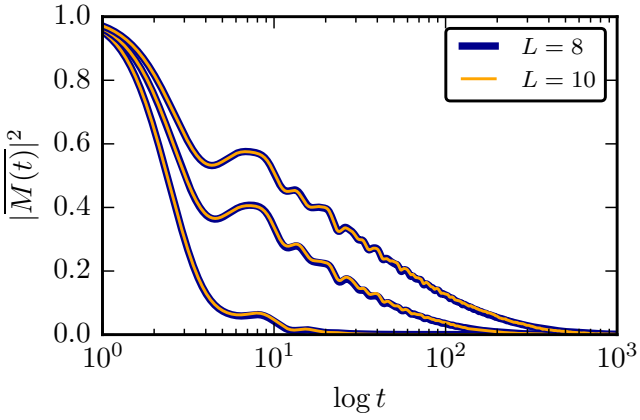

Figure 5.7: The curves for $|\overline{M(t)}|^2$ for system sizes $L = 8$ and $L = 10$, for disorder strengths $h_{\max} = 1$, 3 and 5 respectively, from bottom to top. There is no change in the behaviour of the decoherence, even on an exponential scale in time. Small deviations are due to an order of magnitude fewer disorder realizations for $L = 10$.

can be measured using spin fluorescence. Another possible realization may lie in the use of cold atomic systems, where the spin degree of freedom is e.g. represented by the presence or absence of an atom on a given site.

IV

Conclusion and Appendices

6

CONCLUSION

In this thesis we have dealt with both equilibrium and non-equilibrium condensed matter physics. In Chapter 2, we have discussed the AKLT state in order to introduce the concepts of symmetry protected topological (SPT) states of matter and matrix product states (MPS). We have restricted ourselves to the particular case of SPTs, although the concept of topological states of matter is much broader than we have discussed. In Chapter 4 then, we have investigated the fate of symmetry topological states of matter in one dimension when coupled to an environment [82]. We concluded that the topologically non-trivial state of matter is lost upon reaching the steady state. It is not clear from our analysis whether this is a general truth that also holds for higher dimensions, or for the other types of topological states of matter. These are fundamental questions that are of importance not just for the understanding of these systems, but also for potential applications. The long-range entangled topological states of matter for example, are known to host non-Abelian quasi-particles as low lying excitations. Their non-commuting properties under exchange operations makes them excellent candidates for topological quantum computation.

Our analysis in Chapter 4 was performed using a numerical algorithm (iTEBD) based on matrix product states, that was extended to work on density matrices instead of the usual pure states. As a research tool by itself, this numerical method may be applied to many other systems. In fact, it has shown to be a most valueable tool in research on driven and dissipative photonic cavity systems [138], which by their very nature require a description as an open system.

The representation in terms of matrix product states also allowed for a useful classification scheme based on the projective representation of symmetry operators. Although the MPS formalism is not necessary for this, it provides an elegant description of the results in a diagrammatic form. We have shown how to apply the classification scheme on

the mixed state MPS, and have shown that it is possible to have genuine mixed states that have non-trivial projective representations.

In Chapter 5 of this dissertation we studied a different kind of out-of-equilibrium system. Instead of considering the coupling to an environment in terms of Lindblad operators as we did in Chapter 4, here we have described the combined dynamics of a spin probe and a system that it couples to. By considering only the time evolution of the probe, we may regard the system as the environment explicitly. The non-equilibrium nature in this project arises due to the initial conditions, which are such that the whole system is out-of-equilibrium and relaxes to it as a function of time. In the transient timescales of this dynamics, we have shown that we may learn about the properties of the system by only observing the decoherence of the probe spin [139].

The system we considered was that of a disordered and interacting spin-chain. This spin-chain has a phase transition between a thermalizing and non-thermalizing phase, where the latter is known to be of a many-body localization (MBL) type. Even though we have shown that the spin probe is capable of clearly signalling the MBL phase, the notion of MBL becomes obsolete when the system has access to a heat-bath for exchanging energy or particles. In a real-world situation, such a heat-bath is always present and one should ask the question on which timescales it becomes important. Our current analysis considers the environment of the probe spin as being an isolated system itself. In future work in this direction, it is interesting to take an MBL system and couple it to heat-baths explicitly to determine such timescales. Experiments in this direction have already shown that the timescale is in princple long enough to observe the MBL physics, but it is also clear that the environment is the limiting factor.

Last, in Chapter 3 we have considered the application of machine learning to condensed matter physics. Especially in the case of disordered systems, where a large amount of data is easily available in the form of different disorder realizations, such machine learning methods may provide a valueable tool for analysis. The same goes for finite temperature phase transitions in statistical mechanics, where at a given temperature one may use e.g. Monte Carlo methods to generate large amounts of configurations. Generally speaking, machine learning approaches excel in situations in which large amounts of data with (non-)obvious patterns are available.

We have used a semi-supervised method to discover the correct labelling for a given set of data [68]. This has allowed us to find phase transitions from the entanglement spectra of wavefunctions. Of course, this method does not rely on the underlying data to be in the form of entanglement spectra. It rather relies on patterns in the data, and on a neural network being able to extract these patterns. We have not performed an extensive comparison to other machine learning approaches that try to extract patterns or provide classifications. What is clear however, is that our method is not the fastest way of feature extraction in all cases. Whenever a linear pattern exists in the data, a principal component analysis provides much faster and much easier-to-interpret results. Our method on the other hand, is more powerful than such linear feature extraction, since a neural network can find any relation in the inputs. It would be interesting to investi-

gate *how* a network manages to learn these patterns in a robust way, possibly through the inclusion of other regularization methods or even other types of neural networks. Other interesting routes would be the consideration of more approaches from artificial intelligence and machine learning, such as for example the use of genetic algorithms to optimize wavefunctions.

A

MATRIX PRODUCT STATES

This appendix is a standalone introduction to matrix product states (MPS), and includes a short explanation of the implementation of a DMRG and TEBD algorithm. The subject of pure state MPS is introduced from various viewpoints, after which we also consider the extension of MPS to simulating mixed states. The vectorization procedure that is required for this is discussed, and the method of turning superoperators into matrices. Lastly, we notice the minor differences in measuring observables in the mixed state case compared to the pure state case, and suggest a possible resolution.

A

A.1. GENERAL

The concept of matrix product states (MPS) can be introduced at different levels of complexity. Some are initially more insightful than others, but require a longer exposition to arrive at a working understanding. If all that is needed is being able to work with MPS representations however, many of the underlying elegant details can be left out.

In the following sections, we will provide such different introductions to MPS representations. First, a straightforward version that immediately justifies the *matrix product* nomenclature and provides the necessary details for implementation. Next, MPS representations are introduced by way of a graphical method based on finite state automata. The latter description allows for a deeper insight into the structure and restrictions of matrix product states, but is further removed from a direct implementation. Finally, a more rigorous introduction highlighting the underlying concepts is presented at the end of this section.

Before we begin, a few general remarks are in order. Matrix product states are a way of representing a quantum state of a one dimensional (1D) system. The logic behind this particular representation extends to higher dimensions, for which it is in general referred to as a tensor network. In other words, MPS are the 1D variant of tensor networks. Once we arrive at a diagrammatic representation of an MPS, such conceptual extension to higher dimensions is straightforward. As a matter of fact, one may choose to start the introduction to MPS by beginning with general tensor networks and then specifying to 1D. It should be noted that the actual implementation of algorithms for tensor networks in dimensions larger than one is much more complicated [140].

To set the stage, we consider a general chain of sites as depicted in Figure A.1. Every site

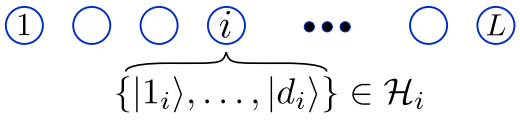

Figure A.1: A 1D system consisting of sites labeled by an index i, each having associated to it a local Hilbert space $\mathscr{H}_i$ of dimension d_i. The states in this Hilbert space are denoted by $|s_i\rangle$, with $s_i = 1 \ldots d_i$.

of the 1D chain is labeled by an index i and has associated to it a local Hilbert space $\mathscr{H}_i$. The dimension of this local Hilbert space is $\dim(\mathscr{H}_i) = d_i$, such that e.g. for a spin-1 chain we have $d = 3$ for all sites. It is possible to mix different Hilbert spaces for different sites, s.t. models that mix different spin-types or even fermions and bosons can be represented. The possible states for each site are labeled $|s_i\rangle$, with $s_i = 1 \ldots d$ (see Fig. A.1), with which we can form basis states for an L site system by constructing product states $|s_1 s_2 \ldots s_L\rangle = |s_1\rangle \otimes |s_2\rangle \otimes \ldots \otimes |s_L\rangle$. An arbitrary pure state wavefunction for this system can be expressed in this basis as

$$|\psi\rangle = \sum_{s_1,\ldots s_L} c_{s_1 \ldots s_L} |s_1 \ldots s_L\rangle. \tag{A.1}$$

Given that the basis states themselves are not time-dependent, the (time-)evolution of this wavefunction can be represented solely by the coefficients $c_{s_1 \ldots s_L}$. The number of coefficients in this wavefunction is d^L, which for a full many-body system quickly becomes impossibly large as a function of the system size. Impossible refers to the possibility of numerically computing this wavefunction. Nowadays, many quantum many-body problems can only be satisfactory dealt with using the help of numerics. This means that the wavefunction has to be stored in memory. As an example consider a chain of two-level systems, s.t. $d = 2$ everywhere. Storing a wavefunction for L sites in memory, requires at the very least 2^L bits (each bit indicating the state of the two level system on a site). For a computer with 64 gigabytes of RAM, this amounts to maximally $L = 40$. It suffices to say that this would already require an optimized algorithm, let alone the required CPU-hours to perform any meaningful calculation with such a wavefunction.

The possibility of the enormous number of coefficients in Eq. (A.1) is not necessarily a problem. Namely, we are usually not after the most general wavefunction of a system. Rather, we tend to target the groundstate and possibly a few other low-energy eigenstates. These states are far from random as far as the coefficients are concerned, and we may wonder if for such states we could do with less than exponentially many coefficients. It is not hard to imagine that low-energy many-body eigenstates have some structure, and that hence many of the coefficients are identical or even zero. We should look for a different way of writing the coefficients that capitalizes on the fact that coefficients are not completely random. Matrix product states turn out to provide this notation, the reasons for which will be made clear in the following sections. To put this into perspective, whereas the exact representation of the general wavefunction in Eq. (A.1) for two-level systems limits us to about $L = 40$, using MPS representations can boost this to 1000+ sites.

A.1.1. HEURISTIC INTRODUCTION

Without loss of generality it is possible to write the general wavefuncion of Eq. A.1 in a MPS form, by substituting

$$c_{s_1 \ldots s_L} \rightarrow \mathrm{tr}\big(M^{[1],s_1} \cdots M^{[L],s_L}\big). \tag{A.2}$$

Notice that we have given a 'local' structure to the coefficients: for every site of the chain we have introduced a set of d matrices. The matrix product between all of the local matrices is what gives the MPS its name. The trace over this product of matrices is there to enforce that the result is a scalar. Also notice that we have not thrown away any coefficients; there are still d^L coefficients on the right-hand side, although we now have a way of controlling the number of *distinct* coefficients. This control is hidden in the dimensions of the matrices M, which therefore plays a very important role. We will see later that this also controls the amount of entanglement the wavefunction can have, which already shows that MPS representations work well for states with 'little' (to be made more precise later, in section A.1.3) entanglement. Letting the dimension of the matrices be $D \times D$, the right-hand side roughly has dLD^2 distinct coefficients. Hence in order to represent a fully random wavefunction in which all of the d^L coefficients are distinct

we would need D to scale exponentially in system size. This means that the substitution in Eq. (A.2) is *not* an approximation, as long as we choose the dimensions of the M-matrices large enough. Numerical algorithms based on MPS representations will of course limit the maximum value for D, turning the MPS representation in an approximation to the actual wavefunction.

The class of states with 'little' entanglement mentioned before, are those that can exactly be represented by MPSs with matrix dimensions that scale at most linear in L. In section A.1.3, we examine this class of states and construct the MPS representation explicitly. For the sake of completeness, we write out a general wavefunction in MPS form:

$$|\psi\rangle = \sum_{s_1,\ldots s_L} \mathrm{tr}\big(M^{[1],s_1} \cdots M^{[L],s_L}\big)|s_1 \ldots s_L\rangle. \tag{A.3}$$

For most practical purposes, it suffices to start from this representation. As a matter of fact, initializing a new MPS can be done in just a few lines as shown in the `Python` code snippet below. Numerical algorithms such as the *density matrix renormalization group* (DMRG) or *time evolving block decimation* (TEBD) both provide a set of rules on how to update the M-matrices to perform a groundstate search or time evolution, respectively. More on these algorithms follows in section A.2

```
import numpy as np
# Set the MPS parameters
d = 2; D = 10; L = 40
# The MPS matrices M are stored here as a set of L tensors,
# each having d matrices of size DxD. We initialize them with zeros,
# but other initializations are also possible.
M = np.zeros( (L,d,D,D ) )
```

A.1.2. FINITE STATE AUTOMATA AND MPS

Let us take a step back and forget about the MPS representation for the moment. Instead, we draw inspiration from the observation that groundstate and low-energy excited states of various many-body systems tend to have a specific structure, or order. This order could be straightforward, as it is in a ferromagnet, but it may very well be more complicated as in the case of non-local order in spin-1 antiferromagnets. For such states, we may ask which of the basis states (written in the local product basis as in Eq. A.3) have the correct order and may therefore participate in its description. As an example, the groundstate of a spin-1/2 ferromagnet is very simple, and only one basis state participates: $|\uparrow\uparrow \ldots \uparrow\rangle$ (or of course all spins down). A similar idea can be used to construct a fully delocalized single spin, in which all basis states of type $|\downarrow \ldots \downarrow\uparrow\downarrow \ldots \downarrow\rangle$ (in which only one spin is in the $|\uparrow\rangle$ state and all others in the $|\downarrow\rangle$ state) participate.

For the states that we described above, we can think of building a *finite state machine* (FSM), also known as a *finite state automaton*, that would accept the allowed basis states [141]. A finite state automaton is an abstract representation of a machine that can be in a finite number of different states, and is able to transition between them. One of the states is typically chosen to be the initial state, and one or multiple states are labeled as final states. The input to the automaton is a string of symbols, which are processed one-by-one (let us say, from right to left). The automaton transitions from the current state to the next if it reads the corresponding symbol for that transition. If, at the end of the input string, the automaton is in one of the final states, the input is said to have been accepted.

This becomes more clear by explicitly going through an example. Let us build the FSM for accepting the basis states that are allowed for the fully delocalized spin. The input simply consists of the basis state, with the symbols being $|\uparrow\rangle$ and $|\downarrow\rangle$. The FSM for this task needs to have only two states:

- **A**: I have not yet seen a $|\uparrow\rangle$
- **B**: I have seen a $|\uparrow\rangle$

The initial state of the FSM is state **A**, and **B** is the only final state. The allowed transitions for the automaton are as follows:

- **A** to **A**: We stay in state **A**, if we read a $|\downarrow\rangle$.
- **A** to **B**: We transition from **A** to **B**, if we read a $|\uparrow\rangle$
- **B** to **B**: We stay in state **B**, if we read a $|\downarrow\rangle$.
- **B** to **A**: Not allowed. If we read a $|\uparrow\rangle$ in state **B** we terminate.

One can easily check that all states such as $|\uparrow\rangle$, $|\downarrow\uparrow\downarrow\ldots\downarrow\rangle$, $|\downarrow\ldots\downarrow\uparrow\rangle$ are accepted (finish in state **B**), whereas any input with *no* $|\uparrow\rangle$ or more than one $|\uparrow\rangle$ will be rejected (in the former case since we never reach state **B**, and in the latter because reading another $|\uparrow\rangle$ in state **B** terminates the machine). If we wish the FSM to only accept inputs of a certain length, we can add the restriction that only a given number of transitions can be made.

The FSM for the delocalized single spin is graphically represented in Figure A.2. The initial state is indicated by an arrow pointing to it from the outside, and final states are usually indicated by double circles.

We can also turn the FSM into an automaton *generating* the allowed states. In this setting, we no longer feed an input to the machine. Rather, we start it in an initial state and allow for a given number of transitions to occur. Each time a transition takes place, we make the machine *output* the symbol associated to the transition. If we start the automaton in state **A**, and make it perform 10 transitions, it can generate *almost* all the allowed basis states for the delocalized spin for $L = 10$. The two boundary states, in which

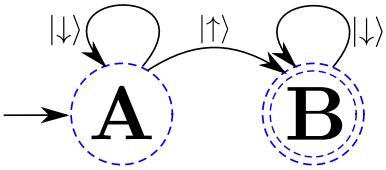

Figure A.2: A simple finite state automaton for accepting/generating a wavefunction for a fully delocalized single spin.

the spin-up is on the first or last site, require a bit more thought. This problem can easily be fixed however, by including the option of starting the machine in a arbitrary state. Like this, when the first site has a spin-down or spin-up we start in the **A** state or in the **B** state respectively.

What may be non-obvious at this point, is that a direct mapping exists between such an FSM and the M matrices in the MPS representation of Eq. A.3. The number of different states in the automaton gives the (minimally required) dimension D of the MPS matrices, while the number of different labels on the transitions is what sets the local Hilbert space. This means that the machine of Figure A.2 is mapped to the following set of matrices:

$$M_{\downarrow} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad M_{\uparrow} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}. \tag{A.4}$$

The $M_{\downarrow}$ matrix is diagonal, since we output a $\downarrow$ only when we transition from **A** to **A** or from **B** to **B**. Similarly, the only non-zero entry in $M_{\uparrow}$ is on the offdiagonal representing the transition from **A** to **B**.

A very important point to make about the FSM approach, is that it does not provide a way of computing the weights of each of the basis states directly (see the next example). In this sense, the FSM approach is severly lacking. It does, however, provide an intuitive explanation for the construction of various matrix product states.

For completeness and in order to provide a reference for later, the FSM in Figure A.3 generates the groundstate wavefunction of the spin-1 AKLT model in the zero-magnetization sector. This groundstate is characterized by non-local string order, meaning that the $s_z = +1$ and $s_z = -1$ values are perfectly alternating with an arbitrary number of $s_z = 0$ in between.

A.1.3. RIGOROUS INTRODUCTION

This time, we are motivated by the observation that groundstates and low-energy excited states of various many-body systems tend to have 'little' entanglement. This is a statement that can be made rigorous [142] for 1D systems that have a gap between the groundstate and the first excited state. It has consequences that we will fully exploit to

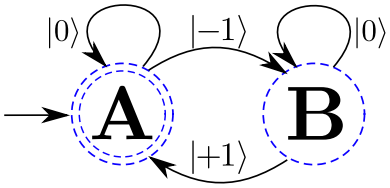

Figure A.3: The automaton generating the allowed basis states for the AKLT groundstate, i.e. those basis states that have string order. There are only two states, indicating that matrix dimension $D = 2$ is enough to represenent such a non-trivial state. Notice that the initial state is also the final state, since we want to select the groundstate in the zero-magnetization sector. The state with only $|0\rangle$ is also generated by this FSM but would have zero weight in the actual groundstate, highlighting the weakness of the FSM approach.

obtain the MPS representation.

First, let us make solid what is meant by entanglement. We consider the so-called *bipartite* entanglement between two parts of a system for any bipartitioning. The two parts are generally refered to as the *left* and *right* half, or as parts A and B, and are associated with two Hilbert spaces $\mathscr{H}^A$ and $\mathscr{H}^B$ of dimensions $\dim \mathscr{H}^{A/B} = d^{A/B}$. These Hilbert spaces depend on where the bipartitioning is made, i.e. at which point the system is separated into two parts. By introducing a basis for parts A and B, denoted $\{|\phi_i^A\rangle\}$ and $\{|\phi_j^B\rangle\}$ respectively, we can write any wavefunction in a productbasis between A and B as follows:

$$|\psi\rangle = \sum_{ij} \alpha_{ij} |\phi_i^A\rangle |\phi_j^B\rangle. \tag{A.5}$$

The beauty of the *Schmidt decomposition*, is that this double sum over the coefficients α_{ij} and basis states can be turned into a single sum

$$|\psi\rangle = \sum_{\alpha=1}^{r} \lambda_\alpha |\Phi_\alpha^A\rangle |\Phi_\alpha^B\rangle, \tag{A.6}$$

where the Schmidt rank r is constrained between 1 and the minimum of the Hilbert space dimensions of parts A and B, i.e. $r = 1 \ldots \min(d^A, d^B)$. The coefficients λ_α are known as the Schmidt coefficients. The Schmidt coefficients are real, non-negative, unique (for a given $|\psi\rangle$) and satisfy $\sum_\alpha \lambda_\alpha^2 = 1$. The Schmidt rank is equal to 1 only for a product state, which by definition is not entangled, whereas a Schmidt rank $r > 1$ indicates non-zero entanglement between the two parts. The basis states $|\Phi_\alpha^A\rangle$ and $|\Phi_\alpha^B\rangle$ for the A and B parts are related to the original $|\phi_i^A\rangle$ and $|\phi_j^B\rangle$ by a unitary transformation that depends on the state itself. Namely, in order to get to Eq. A.6 from Eq. A.5 we perform a singular value decomposition (SVD) of the coefficient-matrix α_{ij}:

$$\alpha_{ij} = \sum_\alpha U_{i\alpha} S_{\alpha\alpha} V_{\alpha j}^\dagger,$$

where S is a diagonal matrix with entries λ_α, and the matrices U and V have the properties $UU^\dagger = 1$ and $V^\dagger V = 1$ (with $\dagger$ representing the hermitian conjugate). The diagonal

entries of S are also called the *singular values*, and give rise to the properties of the λ_α coefficients mentioned before.

As a side note, we consider a different route of obtaining the Schmidt coefficients and the vectors $|\Phi_\alpha^{A/B}\rangle$. Starting from the density matrix $\rho = |\psi\rangle\langle\psi|$ of the combined parts A and B, we can obtain both the coefficients and vectors from performing a partial trace of ρ:

$$\rho_A = \mathrm{Tr}_B \rho = \sum_\alpha \lambda_\alpha^2 |\Phi_\alpha^A\rangle\langle\Phi_\alpha^A|$$

and similar for ρ_B. The two reduced density matrices ρ_A and ρ_B share the same spectrum λ_α^2, and the eigenvectors form the orthogonal columns of U and $V^\dagger$.

From the Schmidt values λ_α, or from the reduced density matrices, we may construct a quantity known as the *entanglement entropy*. The entanglement entropy is nothing more than just the regular von Neumann entropy, but evaluated on a subsystem (namely, the reduced density matrix):

$$S = -\mathrm{Tr}\rho_{A/B} \ln \rho_{A/B} = -\sum_\alpha \lambda_\alpha \ln \lambda_\alpha. \tag{A.7}$$

Only for a product state is $S = 0$, hence any nonzero S signals the presence of entanglement [1]. As a remark for later, we should note that the Schmidt values λ_α are also referred to as the entanglement spectrum. Rather, to be precise, it is $-2\ln\lambda_\alpha$ that constitutes the entanglement spectrum, but the term is used interchangeably for both.

We can now come back to explaining what is meant by 'little' entanglement. Consider the entanglement entropy as a function of the size N of the left half of the bipartition, so that $N = 1 \ldots L/2$. We need only go up to half the system size, since anything beyond that is equivalent to having made the bipartitioning starting from the opposite end. Let us denote this quantity $S(N)$, and ask how it scales as a function of N. This will of course depend on the particular state under consideration, but it is possible to prove that $S(N)$ obeys an *area law* for groundstates of gapped 1D systems. The area law states, independent of dimension, that $S(N)$ is proportional to the size of the boundary that is left after a bipartioning. In 1D systems, the boundary consists of 2 points and hence the entanglement entropy does not scale with N. Near the boundaries of a system there can of course be a dependence on N (boundary effects), but the entanglement entropy quickly satures as is demonstrated in Figure A.4.

Let us now make use of the area law, by constructing *explicity* such a state for which $S(N)$ is constant. As a first step, we imagine a chain of sites as depicted in Figure A.5. Each of the sites is associated with a complex Hilbert space $\mathbb{C}^D$ for some positive $D \geq 1$, and we imagine pairing the sites into groups of $\mathbb{C}^D \times \mathbb{C}^D$ as shown in the figure. We think of these pairs as *one* new site. The pairs are sometimes referred to as auxiliary sites. Next, we construct *bonds* between the pairs, such that for example the state of auxiliary sites 2

[1] Actually, the entanglement entropy also captures classical correlations and may therefore not be the best entanglement 'witness'. Other quantities, such as the logarithmic negativity [143] have been proposed, but finding an optimal witness is a research project by itself.

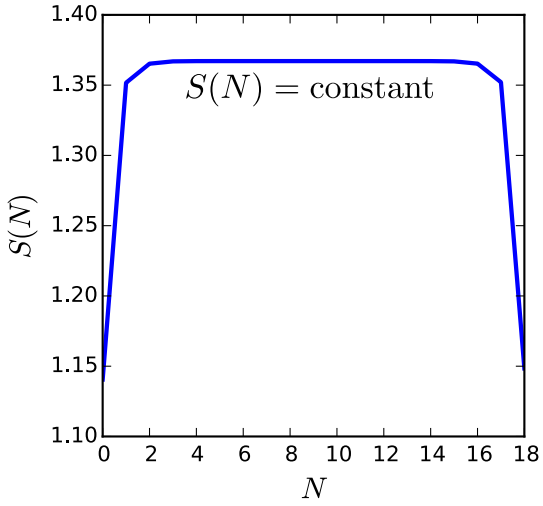

Figure A.4: The entanglement entropy as a function of bipartition size N for the groundstate of a gapped 1D system with $L = 20$ (the AKLT state, see section 2.2) The sites are indexed with labels starting from 0. Apart from the effects at the boundary, it is clearly visible that the entanglement entropy saturates and becomes independent of N (area law).

and 3 is given by

$$|\omega_D^{23}\rangle = \frac{1}{\sqrt{D}} \sum_{k=1}^{D} |k_2\rangle|k_3\rangle, \tag{A.8}$$

where $|k_2\rangle$ represents the state of site 2 with $k = 1 \ldots D$. The quantity D is referred to as the (maximal) *bond dimension* for the MPS. The full state for the chain is then given by

$$|\psi\rangle = |\omega_D\rangle^{\otimes L-1}, \tag{A.9}$$

where we construct a bond between all of the pairs.

We will be considering open boundary conditions mostly, and so site number 1 and $2L$ are left unpaired. For periodic boundary conditions, an extra bond between those two sites has to be added. Each of the bonds corresponds to a maximally entangled state, as we can demonstrate on the example bond $|\omega_D^{23}\rangle$. The density matrix for this bond is given by

$$\rho^{23} = |\omega_D^{23}\rangle\langle\omega_D^{23}| = \frac{1}{D} \sum_{k,l=1}^{D} |k_2\rangle\langle l_2| \otimes |k_3\rangle\langle l_3|, \tag{A.10}$$

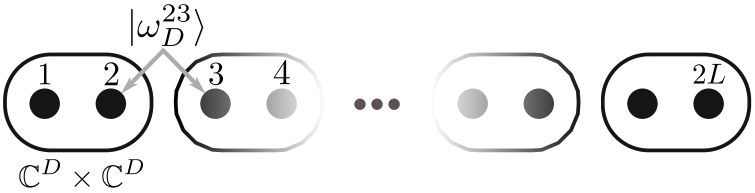

Figure A.5: Construction of maximally entangled bonds $|\omega_D\rangle$ between 'sites'. Each site consists of two auxiliary sites and is assigned to a space $\mathbb{C}^D \times \mathbb{C}^D$, so that a chain with L sites actually consists of $2L$ auxiliary sites.

leading to a reduced density matrix

$$\begin{aligned}\rho^{(2)} &= \mathrm{Tr}_3 \rho^{23} \\ &= \frac{1}{D}\sum_{\alpha=1}^{D}\sum_{k,l=1}^{D} |k_2\rangle\langle l_2| \underbrace{\langle\alpha_3|k_3\rangle}_{=\delta_{\alpha,k}} \underbrace{\langle l_3|\alpha_3\rangle}_{=\delta_{l,\alpha}}, \\ &= \frac{1}{D}\sum_{\alpha=1}^{D} |\alpha_2\rangle\langle\alpha_2|. \end{aligned} \tag{A.11}$$

The entanglement entropy is then easily found to be

$$\begin{aligned} S^{23} &= -\mathrm{Tr}(\rho^{(2)} \ln \rho^{(2)}) \\ &= -\sum_k \frac{1}{D} \ln \frac{1}{D} \\ &= \ln D, \end{aligned} \tag{A.12}$$

and hence for any bipartitioning, the state $|\psi\rangle$ in Eq. (A.9) has entropy equal to $\ln D$. Notice here that we do not consider bipartitionings in between the bonds. The two sites in a pair are regarded as *one* physical site that cannot be cut in half. It is clear here that the bond dimension sets the maximal entanglement entropy that can exist across any bond. Apart from the constant entanglement entropy, the current state is actually not very interesting. To turn it into an interesting state, we introduce maps $P^{[i]}$ from the pairs in $\mathbb{C}^D \times \mathbb{C}^D$ to the 'physical' Hilbert space $\mathbb{C}^d \sim \mathscr{H}_i$, as shown in Fig. A.6. We obtain a new state by pairwise application of these maps:

$$|\tilde{\psi}\rangle = P^{[1]} \otimes \cdots \otimes P^{[N]} |\psi\rangle = P^{[1]} \otimes \cdots \otimes P^{[N]} |\omega_D\rangle^{\otimes L-1}$$

From the field of quantum information theory it is known that local operators (such as our P) do not change the entanglement. For an operator to change the entanglement between two parts of a bipartition, it needs to act on both parts simultaneously. Our P-maps never act across a bond that can be cut. In other words, we still have a state that is maximally entangled across each bond, but we can add structure (correlations) to the state via our P-maps.

Let us try to write an explicit form for the P-maps. The basis for the local Hilbert spaces is denoted $|s\rangle$ as before, but we additionally introduce basis states $|\alpha\beta\rangle$ for the pairs living

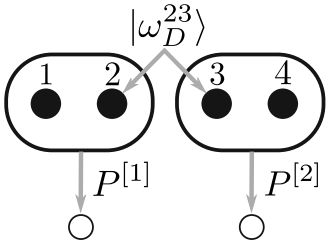

Figure A.6: We construct maps $P^{[i]}:\mathbb{C}^D\times\mathbb{C}^D\to\mathbb{C}^d$, mapping the auxiliary sites to physical sites. These maps map the auxiliary wavefunction in Eq. A.9 to a physical wavefunction that has more interesting correlations, but is still a maximally entangled state across each bond.

in $\mathbb{C}^D\times\mathbb{C}^D$ (so that α and β take values $1\dots D$). For a given P mapping to site i, we then write

$$P^{[i]}=\sum_{\alpha,\beta=1}^{D}\sum_{s=1}^{d}|i\rangle\langle\alpha\beta|\,M_{\alpha,\beta}^{[i],s}, \tag{A.13}$$

in which the coefficients are written compactly via the tensors $M_{\alpha,\beta}^{[i],s}$. Hence for a given site i we have d matrices $\left(M^{[i],1},\dots,M^{[i],d}\right)$, and we associate $P^{[i]}$ to this set of matrices $M^{[i]}$. Notice in particular that we only have to specify ND^2d parameters to completely determine all of the P-maps. Next, instead of a single site, let us look at a two-site case explicitly. We consider the situation shown in Fig. A.6, in which we have two physical sites that come from four auxiliary sites connected by a maximally entangled bond $|\omega_D^{23}\rangle$. We act with the P-maps:

$$\begin{aligned}P^{[1]}\otimes P^{[2]}|\omega_D^{23}\rangle&=\sum_{s_1,\alpha,\beta}\sum_{s_2,\gamma,\delta}|s_1s_2\rangle\langle\alpha\beta\gamma\delta|\,M_{\alpha,\beta}^{[1],s_1}M_{\gamma,\delta}^{[2],s_2}\sum_k\frac{1}{\sqrt{D}}|k\rangle_2|k\rangle_3\\&=\sum_{s_1,s_2,\alpha,\beta,\gamma,\delta}\frac{1}{\sqrt{D}}M_{\alpha,\beta}^{[1],s_1}M_{\gamma,\delta}^{[2],s_2}\,|s_1s_2\rangle\langle\alpha\delta|\underbrace{\langle\beta\gamma|kk\rangle}_{\delta_{\beta,k}\delta_{\gamma,k}}\\&=\sum_{s_1,s_2,\alpha,\delta,k}M_{\alpha,k}^{[1],s_1}M_{k,\delta}^{[2],s_2}\,|s_1s_2\rangle\langle\alpha\delta|.\end{aligned} \tag{A.14}$$

In the last line, we absorbed the $1/\sqrt{D}$ factor into the M-matrices. What we learn from this example is that when the matrices $M^{[1],s_1}$ and $M^{[2],s_2}$ are associated to $P^{[1]}$ and $P^{[2]}$ separately, we can associate the matrix product $M^{[1],s_1}\cdot M^{[2],s_2}$ to the combination $P^{[12]}=P^{[1]}\otimes P^{[2]}$. By iterating this procedure, we will arrive at

$$|\tilde{\psi}\rangle=\sum_{s_1,\dots,s_L,\alpha,\delta}M^{[1],s_1}\cdot\ldots\cdot M^{[L],s_L}|s_1\dots s_L\rangle\langle\alpha\delta|. \tag{A.15}$$

This expression is basically the MPS representation of Eq. A.3, except for the $\langle\alpha\delta|$ term. By inspecting Eq. A.14, we see that the indices α and δ belong to the first and last M-matrices, respectively. They therefore determine the boundary conditions on the MPS, and choosing them is an important part of MPS representations.

To get to a fully physical wavefunction, we need to eliminate the $\langle\alpha\delta|$ term. There are two main ways of achieving this, keeping in mind that the final product of the M-matrices should result in a scalar (they are, after all, the coefficients to the basis states).

The first method, often chosen for *open* boundary conditions, is to take the inner product of Eq. A.15 with $|1,1\rangle$ on the first and last auxiliary indices. This leads to fixing $\alpha = \delta = 1$, and thereby turns the first and last matrix into a row and column vector respectively. Since the coefficients themselves have no idea about the boundary conditions, it is not entirely true that this choice is valid only for open boundary conditions. For periodic boundary conditions the restriction of α and δ to 1 will also work, but the entanglement across the periodic bond would have to be encoded throughout all of the other M-matrices. This most likely means that the bond dimension D has to be a (very) large number. Instead, it is more efficient for periodic boundary conditions to keep the first and last M as matrices instead of vectors, and turn the matrix product into a scalar by taking the trace. Instead of multiplying by $|1,1\rangle$, we multiply by $|\gamma,\gamma\rangle$ and sum over γ. This procedure will lead to exactly the formulation of an MPS state as in Eq. A.3. For notational convenience, we will always write the MPS using the trace form, even for open boundary conditions in which we turn the first and last M into row and column vectors. The trace will then simply act on a scalar and change nothing. In the following section, we will introduce the diagrammatic way of working with MPS representations. It is here where the difference between periodic and open boundaries can be clearly seen.

As a final note, there also exist infinite MPS (iMPS) representations. In this case, we consider a unit cell of L sites, and impose periodic boundary conditions. A crucial difference to a finite periodic case, is the way in which observables are evaluated. More on this will be highlighted in the following sections, but it is good to be aware of the existence of iMPS already now.

A.1.4. MPS DIAGRAMS

Formulating operations on MPS states can be done very efficiently by introducing a diagrammatic representation for the M-matrices. The simplest object we consider is exactly such a single M-matrix, and is shown below in Fig. A.7. The lines in this diagram indicate the indices of the M-tensor, also referred to as its 'legs'. The diagram hence has as many legs as the object it represents has indices. The downward pointing leg is called the physical leg, since it contains the indices belonging to the physical basis states. The other two legs, simply called the left and right legs, represent the auxiliary indices running from 1 up to D. With this representation, we can re-express common operations on

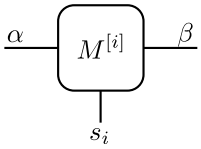

Figure A.7: MPS diagram for a single M-tensor. The number of 'legs' is the number of indices of the tensor, which is three for a single M-tensor. The downward pointing leg (convention) is chosen to be the physical leg, having the index corresponding to the physical Hilbert space. The left and right legs are auxiliary legs.

the tensors in a graphical way. The most basic and most important operation one can perform is a *contraction* of two tensors over one or multiple of their indices. An example of this is shown in Fig. A.8. A contraction of two tensors over one or multiple legs can

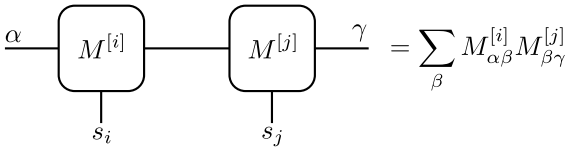

Figure A.8: Contracting two (or more) legs of tensors is a graphical way of representing a summation over the corresponding indices of the tensors. Only legs with the same dimension can be contracted over. In this example, the resulting object is a tensor with 4 open indices, and hence would be drawn as an object with 4 legs.

be implemented straightforwardly in `Python`. The code snippet below and performs the contraction shown in Fig. A.8, albeit with tensors that have been initialized to zero.

```
# ...
M = np.zeros( (L,d,D,D ) )
# Contract MPS matrices of site 4 and 5, using the 'right' leg of
# site 4 (index 2) and the 'left' leg of site 5 (index 1). Notice
# that the resulting object has four legs, the first index being
# the physical leg of site 4, the second being the left leg of site
# 4, then the physical leg and right leg of site 5.
newTensor = np.tensordot( M[3], M[4], axes=(2,1) )
```

A general MPS as in Eq. A.3, reproduced below for convenience, can now be represented diagrammatically. As alluded to in the previous section, it is here where we can distinguish clearly between open and periodic boundary conditions. As is shown in Fig. A.9, for open boundary conditions the first and last M-matrices are vectors and hence have one index less than do the matrices[2]. In the case of periodic boundaries, the trace amounts

[2] In principle, we could use the same diagram also for the vectors, but restricting the legs to running only over one value. But that means there is nothing to contract anyway, so we may as well drop the leg.

to a sum over the α and δ indices in Eq. A.15 and is represented by the line connecting the first and last MPS matrix.

$$|\psi\rangle = \sum_{s_1,\dots s_L} \mathrm{tr}\big(M^{[1],s_1}\cdots M^{[L],s_L}\big)|s_1\dots s_L\rangle.$$

Notice that some of the labels have been left out on purpose, but will be put back whenever more detail is needed. We will not consider periodic boundaries on MPS in the rest

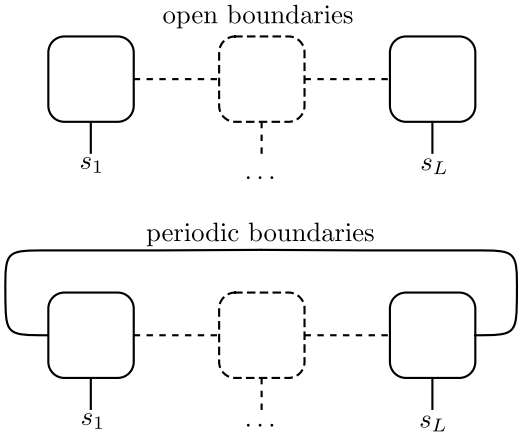

Figure A.9: A general MPS wavefunction is determined fully by the coefficients in a given basis. These coefficients can be represented in an MPS form, for which the diagram is shown here. Hence we think of this diagram as representing the wavefunction (in a corresponding basis). The top figure shows the diagram for an MPS with open boundary conditions, whereas the periodic connection in the lower figure indicates periodic boundary conditions.

of this section, so the diagrams will from now on only consist of those with open boundaries.

The overlap between two wavefunctions $|\psi\rangle$ and $|\phi\rangle$, with MPS matrices denoted by A and B respectively, can then be represented as shown in Fig. A.10. The MPS for $|\phi\rangle$ has to be conjugated (to form $\langle\phi|$), and is drawn with physical indices pointing up. Let us directly turn to the question of how to measure observables. We consider the *transfer matrix*, $\mathbb{E}^{[i]}$, shown in Fig. A.11. It maps auxiliary indices $(\alpha,\alpha') \to (\beta,\beta')$, and can be expressed in terms of the MPS matrices as

$$\begin{aligned}\mathbb{E}^{[i]}_{(\alpha\alpha'),(\beta\beta')} &= \sum_{s_i} M^{[i],s_i}_{\alpha\beta}\left(M^{[i],s_i}_{\alpha'\beta'}\right)^* \\ &= \sum_{s_i} M^{[i],s_i}\otimes\left(M^{[i],s_i}\right)^*. \end{aligned} \tag{A.16}$$

With this notation, we can express the normalization of a wavefunction $|\psi\rangle$ as

$$\langle\psi|\psi\rangle = \mathrm{tr}\Big(\mathbb{E}^{[1]}\dots\mathbb{E}^{[L]}\Big),$$

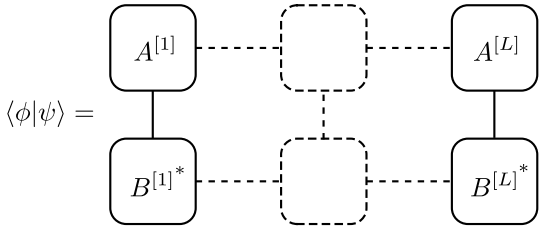

Figure A.10: Overlap of two MPSs $|\psi\rangle$ and $|\phi\rangle$, with matrices labeled by A and B respectively. Notice that the B matrices have their physical indices pointing *upward*, meaning that their matrix entries have to be complex conjugated.

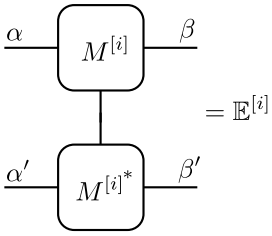

Figure A.11: Mapping auxiliary indices from $(\alpha, \alpha') \to (\beta, \beta')$ by using the transfer matrix $\mathbb{E}^{[i]}$.

where we understand that $\mathbb{E}^{[1]}$ and $\mathbb{E}^{[L]}$ for the first and last sites are slightly different from the transfer matrix shown in Fig. A.11 in that they have only one set of auxiliary indices. That also means that the trace is redundant, but we keep the notation compatible with periodic boundaries. By a slight extension of the transfer matrix, we also introduce $\mathbb{E}^{[i]}_X$:

$$\mathbb{E}^{[i]}_X = \sum_{s_1 s_1'} M^{[i],s_1} X_{s_1,s_1'} \left(M^{[i],s_1'}\right)^*, \tag{A.17}$$

so that $\mathbb{E}^{[i]} = \mathbb{E}^{[i]}_{\hat{1}}$, with $\hat{1}$ the identity operator. We are then capable of expressing an expectation value such as:

$$\langle\psi|X_i \otimes Y_j|\psi\rangle = \mathrm{tr}\left(\mathbb{E}^{[1]} \cdots \mathbb{E}^{[i]}_X \cdots \mathbb{E}^{[j]}_Y \cdots \mathbb{E}^{[L]}\right). \tag{A.18}$$

In a diagram, such an expectation value would look like that shown in Fig. A.12. This is a good point to take a step back and evaluate the efficiency of the computation of an expectation value as in Eq. A.18. When contracting such a full network, it pays to consider attention to the order in which the contractions are carried out. When we perform the contraction shown in Fig. A.13, we have the option of first contracting the pair of legs **A** and then **B**, or vice versa. If we contract the pair of legs indicated by the letter **A** first, we effectively perform multiplication of two matrices of sizes $D^2 \times d$ and $d \times D^2$. Since the complexity of matrix mulitplication of matrices of shapes $n \times m$ and $m \times p$ is $\mathcal{O}(nmp)$, this route costs $\mathcal{O}(dD^4)$. Via the outher route, by contracting the **B** pair first, we multiply two matrices of size $dD \times D$ and $D \times dD$, leading to a complexity $\mathcal{O}(d^2D^3)$. In the limit of very large D, the second route is much better than the first.

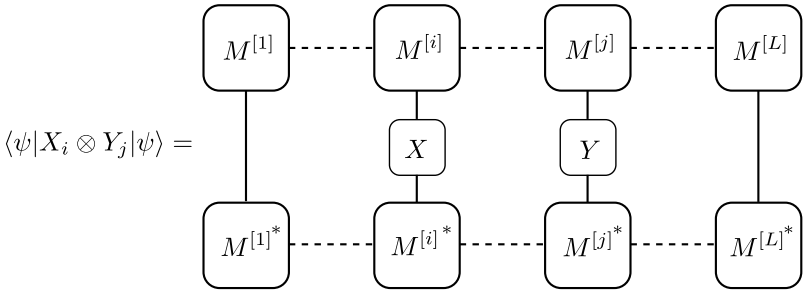

Figure A.12: Computing the expectation value of the operator $X_i \otimes Y_j$ represented in diagrammatical form.

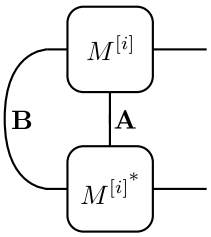

Figure A.13: The order of contracting, i.e. first the pair of legs indicated by **A** or first the pair indicated by **B**, can make a big difference in efficiency. If one contracts **A** first and then **B**, the cost of the full contraction scales as $\mathcal{O}(dD^4)$, whereas the opposite order only costs $\mathcal{O}(d^2D^3)$.

CANONICAL FORM

There is a much more important point that has to be made when we talk about measuring operators. If we want to measure the expectation value of a local operator on a system of length L, we would have to perform the full contraction of all tensors as in Fig. A.12. Instead, there is a degree of freedom in the choice of the MPS matrices that allows us to evaluate local operators by only using the MPS matrices for the site we consider. This degree of freedom is sometimes called a 'gauge' degree of freedom, and amounts to the following transformation on the M-matrices:

$$M^{[i],s_i} \rightarrow X_i M^{[i],s_i} \left(X_{i+1}\right)^{-1}.$$

The matrices X_i and X_{i+1} are unitary matrices, and one can easily convince oneself – by inserting the transformation into Eq. A.3 – that this transformation leaves the MPS unchanged.

A particular choice of the unitary matrices X is given by the 'isometric gauge'. Before going into detail on how to explicitly get the matrices in this form, let us highlight what it will do for us. When chosen correctly, this gauge will bring our M matrices into *canonical*

form. In matrix notation, this means we can choose the MPS matrices to obey

$$\sum_{s_i,\gamma} M^{[i],s_i}_{\alpha\gamma}\left(M^{[i],s_i}_{\gamma\beta}\right)^* = \lambda\delta_{\alpha\beta},, \qquad |\lambda| = 1.. \tag{A.19}$$

Equivalently, canonical form means that $\mathbb{E}^{[i]}$ has as its dominant left eigenmatrix the identity matrix with unit-modulus eigenvalue, i.e.:

$$\sum_{\beta,\beta'} \mathbb{E}^{[i]}_{(\alpha\alpha'),(\beta\beta')}\delta_{\alpha\beta} = \lambda\delta_{\alpha'\beta'}, \qquad |\lambda| = 1. \tag{A.20}$$

If the above conditions are fulfilled, the MPS is called left-canonical. Similary, by straightforward comparison to the left-canonical case, an MPS can also be in a right-canonical form. Diagrammatically, the left-canonical form can be represented as shown in Fig. A.14. Now imagine we are trying to compute the expectation value of an operator X on site i,

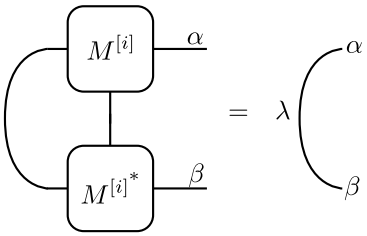

Figure A.14: Left-canonical form means that the contraction of the two MPS matrices over their physical and left indices can be reduced, effectively eliminating the need to actually perform the contraction numerically.

denoted as $\langle\psi|X_i|\psi\rangle$. By making sure that all the MPS matrices to the left of site i are in left-canonical form, and those to the right are all in right-canonical form, we are able to measure a local observable by contracting only three objects as shown in Fig. A.15. This is quite a simplification over a diagram such as the one in Fig. A.12. Practically, let

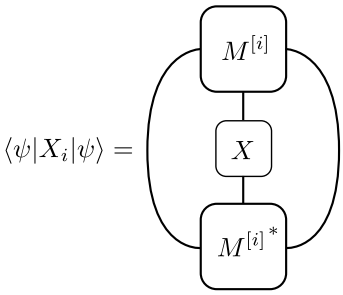

Figure A.15: When the MPS is in canonical form, we no longer need to contract all L sites but can instead focus on just the sites involved in the operator we are measuring.

us look at how the left- or right-canonical form can be obtained, starting from a given

set of MPS matrices M. A canonical MPS can be obtained directly from the coefficients $c_{s_1 \ldots s_L}$ in any basis too, but the details of this are left to the extensive MPS review by U. Schollwöck [144]. For convenience and for fixing notation for the rest of this section, we write here a general MPS with all the indices shown:

$$|\psi\rangle = \sum_{s_1,\ldots,s_L} \sum_{\alpha_1,\ldots} M^{[1],s_1}_{1,\alpha_1} M^{[2],s_2}_{\alpha_1,\alpha_2} M^{[3],s_3}_{\alpha_2,\alpha_3} \ldots |s_1 \ldots s_L\rangle. \tag{A.21}$$

Notice that the first index of the first M-matrix is fixed to taking only one single value, reflecting our use of open boundary conditions. In order for us to construct a left-canonical MPS, we start from the first site and consider $M^{[1],s_1}_{1,\alpha_1}$. We reshape the left and physical indices of this tensor into a single index, and denote the resulting tensor $M^{[1]}_{(s_1,1),\alpha_1}$. This tensor is now a matrix (it has only two indices), and can be decomposed using a singular value decomposition (SVD):

$$M^{[1]}_{(s_1,1),\alpha_1} = \sum_{\sigma_1} U^{[1]}_{(s_1,1),\sigma_1} S_{\sigma_1,\sigma_1} V^{\dagger}_{\sigma_1,\alpha_1},$$

for which S is a diagonal matrix containing the singular values, and U and $V^{\dagger}$ are unitary matrices obeying $UU^{\dagger} = 1$ and $V^{\dagger}V = 1$. We insert this decomposition into Eq. A.21 and perform the sum over α_1 to obtain a new effective MPS matrix $\tilde{M}$ for site 2:

$$\begin{aligned}
|\psi\rangle &= \sum_{s_1,\ldots,s_L} \sum_{\alpha_1,\ldots} M^{[1],s_1}_{1,\alpha_1} M^{[2],s_2}_{\alpha_1,\alpha_2} M^{[3],s_3}_{\alpha_2,\alpha_3} \ldots |s_1 \ldots s_L\rangle \\
&= \sum_{s_1,\ldots,s_L} \sum_{\alpha_2,\ldots} \sum_{\sigma_1} U^{[1]}_{(s_1,1),\sigma_1} \Big(\sum_{\alpha_1} S_{\sigma_1,\sigma_1} V^{\dagger}_{\sigma_1,\alpha_1} M^{[2],s_2}_{\alpha_1,\alpha_2} \Big) M^{[3],s_3}_{\alpha_2,\alpha_3} \ldots |s_1 \ldots s_L\rangle \\
&= \sum_{s_1,\ldots,s_L} \sum_{\alpha_2,\ldots} \sum_{\sigma_1} A^{[1],s_1}_{1,\sigma_1} \tilde{M}^{[2],s_2}_{\sigma_1,\alpha_2} M^{[3],s_3}_{\alpha_2,\alpha_3} \ldots |s_1 \ldots s_L\rangle.
\end{aligned}$$

Note that we have also renamed and reshaped $U^{[1]}_{(s_1,1),\sigma_1} \to A^{[1],s_1}_{1,\sigma_1}$ in the last line. This automatically guarantees that $A^{[1]}$ is left-canonical, c.f. Eq. A.19. This procedure can now be repeated for the next matrix, $\tilde{M}^{[2]}$. It is first reshaped into a matrix, by grouping the left and physical indices: $\tilde{M}^{[2],s_2}_{\sigma_1,\alpha_2} \to \tilde{M}^{[2]}_{(s_2,\sigma_1),\alpha_2}$. Next, we perform a SVD of this matrix, multiply $S_{\sigma_2,\sigma_2} V^{\dagger}_{\sigma_2,\alpha_2}$ onto $M^{[3]}$ to obtain $\tilde{M}^{[3]}$ and reshape $U^{[2]}_{(s_2,\sigma_1),\sigma_2}$ back into $A^{[2],s_2}_{\sigma_1,\sigma_2}$. Iterating this procedure until the very last site brings the full MPS into left-canonical form. When SVD'ing the last site, the $S_{1,1} V^{\dagger}_{1,1}$ is simply a scalar since $\tilde{M}^{[L]}$ is a column vector. It determines the normalization of the wavefunction, and hence should be equal to 1. The reshaping of the tensors and the SVD can also be done easily in `Python`, shown below in the code snippet.

```
# ...
# We are going to SVD the fourth MPS tensor, getting it into
# a left-canonical form. It is important here that the
# legs to be reshaped are neighboring indices.
reshapedM = np.reshape( M[3], (d*D, D) )
# We force the SVD to keep the full matrices for simplicity here.
```

```
# In an actual implementation, truncating the matrices is a
# beneficial thing to do.
U,S,Vdag  = np.linalg.svd( reshapedM, full_matrices=1 )
# The new M[3] is now set to be the reshaped U matrix, auto-
# matically being in left-canonical form.
M[3]      = np.reshape( U, (d,D,D) )
# We should now multiply S and Vdag, and contract the result
# with tensor M[4]. Then we repeat these steps.
```

If we want to generate a right-canonical MPS, we start from the last site and regroup the *right* index with the physical index before performing the singular value decomposition. The new matrices are then formed from the $V^\dagger$ matrices of the SVD, automatically guaranteeing the right-canonical form.

Before discussing algorithms that can be based on MPS representations, it is worth introducing a slightly different notation for the MPS matrices that facilitates working with the left- and right-canonical form. In this new notation, the MPS is written as

$$|\psi\rangle = \sum_{s_1,\ldots,s_L} \Gamma^{[1],s_1}\Lambda^{[1]}\Gamma^{[2],s_2}\Lambda^{[2]}\ldots\Gamma^{[L-1],s_{L-1}}\Lambda^{[L-1]}\Gamma^{[L],s_L}|s_1\ldots s_L\rangle. \tag{A.22}$$

The $\Lambda^{[i]}$ matrices are diagonal with entries corresponding to the Schmidt values for a bipartitioning after site i. These are of course exactly the singular values obtained from the SVD of matrix $M^{[i]}$. In fact, the easiest way of obtaining this representation is by first constructing a fully left-canonical MPS, keeping track of the singular values and storing them as $\Lambda^{[i]}$. Then we may insert the identities $\Lambda^{[i]}\left(\Lambda^{[i]}\right)^{-1}$ in between the $A^{[i]}$ and $A^{[i+1]}$ matrices, absorbing the inverse matrix into $A^{[i+1]}$ and renaming it $\Gamma^{[i+1]}$. Notice though, that we have left out the final $\Lambda^{[L]}$. This is because, as mentioned before, the last tensor is a column vector and hence $\Lambda^{[L]}$ is simply a scalar. For periodic boundary conditions this would not be the case and we benefit from including a $\Lambda[L]$ matrix.

The elegancy of this notation is that we can easily convert it to left- or right-canonical form as follows. We have already seen, by construction, that the left-canonical matrices can be formed via $A^{[i],s_i} = \Gamma^{[i],s_i}\Lambda^{[i]}$. In order to construct right-canonical matrices, we simply have to contract in the opposite direction: $B^{[i],s_i} = \Lambda^{[i-1]}\Gamma^{[i],s_i}$.

TRUNCATION OF AN MPS

So far, all of the above discussion has been about matrix product states with a general bond dimension D. We have also seen that groundstates of gapped 1D Hamiltonians are efficiently represented by MPS states, meaning that the bond dimension needed does not scale exponentially with the system size (area law). In practice, however, even though the bond dimension has a favourable scaling we will most likely want to limit it to a maximum value. For a given wavefunction that would require a bond dimension D to

be represented exactly, we therefore ask how big of an error we make by limiting the allowed bond dimension to $\chi < D$. The answer is given by the singular values obtained from bringing the state into canonical form. If the MPS is in canonical form, the singular values correspond to the Schmidt coefficients, which are empirically known to fall of exponentially in amplitude for groundstate wavefunctions. Hence truncating the bond dimension can be done in a controlled fashion, by e.g. deciding on a threshold below which the singular values are discarded. A typical threshold number is 10^{-8}, although even smaller threshold values may be taken as well. The truncation error is defined by summing all of the discarded singular values:

$$\epsilon_{\text{trunc}} = \sum_{\lambda_i < 10^{-8}} \lambda_i.$$

If one is more flexible in terms of time and memory requirements of numerical simulations, an alternative route for simulations is fixing the maximal truncation error and adjusting the bond dimension to reach this goal.

INFINITE MPS

There is a variant of matrix product states that will allow us to simulate systems directly in the thermodynamic limit of infinite system size. For this we have to assume translational invariance by introducing a unit cell of size L. For simplicity and for use in the (i)TEBD algorithm, we choose $L = 2$. The MPS wavefunction then reads

$$|\psi\rangle = \sum_{s_1,s_2} \text{Tr}\Big(\Gamma^{[1],s_1}\Lambda^{[1]}\Gamma^{[2],s_2}\Lambda^{[2]}\Big)|s_1 s_2\rangle.$$

The difference to a 2-site periodic system lies in the canonical form. If we choose the MPS matrices such that Eq. A.20 is fulfilled, measurements performed on this system can be thought of as having been contracted 'from infinity'. Seen in the opposite way, the fact that the dominant eigenvalue of the transfer matrix is unity means we could add it as many times as we want to construct a larger and larger system.

A.2. MPS ALGORITHMS

A.2.1. DMRG

The density matrix renormalization group (DMRG) algorithm was initially not formulated using matrix product states. It lends itself very well for such a formulation however. In this section we will shortly highlight the essence of the DMRG algorithm, but will not go into details of how the actual underlying implementation works. The main algorithm that has been used for the projects described in this thesis is the (i)TEBD algorithm, that will be discussed in more detail next.

The goal of the DMRG algorithm is to find the MPS with the lowest energy for a given local Hamiltonian $H = \sum_i h_{i,i+1}$. Each of the operators $h_{i,i+1}$ is a two-site operator. The

algorithm can be extended easily to operators involving more than two sites, at the cost of computational complexity. As a three-phase process, DMRG can be summarized as follows:

1. Start with a randomly initialized (but normalized) MPS wavefunction

$$|\psi\rangle = |\psi[M^{[1]}, M^{[2]}, \ldots, M^{[L]}]\rangle$$

   In this notation, we emphasise that the wavefunction is a function of the M-matrices.

2. We choose a site j, and replace the MPS matrix on that site by a new one labeled X:

$$|\psi[X]\rangle = |\psi[M^{[1]}, \ldots, M^{[j-1]}, X, M^{[j+1]}, \ldots M^{[L]}]\rangle.$$

   We then evaluate the energy of this state:

$$E = \frac{\langle\psi[X]|H|\psi[X]\rangle}{\langle\psi[X]|\psi[X]\rangle}.$$

   At this point, the DMRG algorithm provides a direct way of getting the matrix X_{min} that minimizes the energy. We then set $M^{[s]} = X_{\text{min}}$.

3. Instead of choosing a random site s, we perform a 'sweep': start at site 1 and optimize the MPS matrix. Then move to site 2 and continue all the way up to site L. Then move back all the way to the first site, and we have completed one sweep. This sweeping is then iterated until the energy has converged.

The main part in this algorithm is evidently the second step, i.e. how does one find X_{min}? Consider first the normalization of $|\psi[X]\rangle$. We make sure the MPS is in canonical form, so that the normalization can be computed from X only:

$$\langle\psi[X]|\psi[X]\rangle = \sum_{s_j\alpha\beta} X^{[j],s_j}_{\alpha\beta}\left(X^{[j],s_j}_{\alpha\beta}\right)^* = \mathbf{x}\cdot\mathbf{x},$$

where $\mathbf{a}\cdot\mathbf{b} = \sum_i a_i b_i^*$. This is a useful shorthand notation for the normalization. Next, we turn to evaluating the Hamiltonian:

$$\langle\psi[X]|H|\psi[X]\rangle = \sum_i \langle\psi[X]|h_{i,i+1}|\psi[X]\rangle.$$

If we have replaced site s by the matrix X, there are four general cases that occur for the evaluation of the Hamiltonian. In the first case, i and $i+1$ are both to the left of site j as depicted in Fig. A.16. We can collect the whole left hand side into an object referred to as the *left environment* of site s, denoted $L_{\alpha\alpha'}$. This contraction can be written out as follows:

$$\text{case 1} = \sum_{s_j,\beta,\alpha,\alpha'} X^{[j],s_j}_{\alpha\beta} L_{\alpha\alpha'}\left(X^{[j],s_j}_{\alpha'\beta}\right)^* = \sum_{s_j,\beta,\alpha,\alpha'} X^{[j],s_j}_{\alpha\beta}\left(\tilde{X}^{[j],s_j}_{\alpha'\beta}\right)^* = \mathbf{x}\cdot\tilde{\mathbf{x}}. \tag{A.23}$$

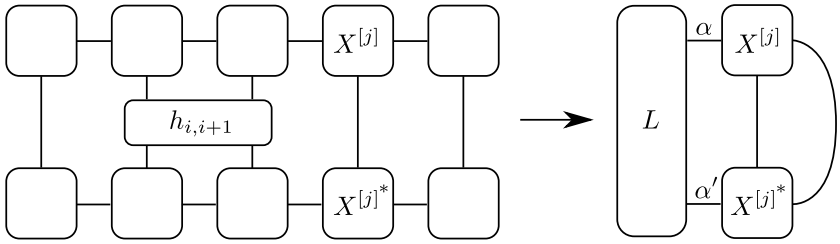

Figure A.16: Definition of the left environment in evaluating the Hamiltonian on an MPS.

This can be further simplified by realizing that $\tilde{\mathbf{x}} = \mathbf{M}\mathbf{x}^*$, in which $\mathbf{M} = \hat{1} \otimes L_{\alpha\alpha'} \otimes \hat{1}$. The first identity operator acts on the physical index of X, whereas the last acts on the right leg. The full expression then turns into

$$\text{case } 1 = \mathbf{x} \cdot \mathbf{M}^* \mathbf{x}.$$

For the other three cases, corresponding to $i = j$, $i+1 = j$ or both i and $i+1$ on the right side of j can be treated similarly. Those cases also lead to a matrix $\mathbf{M}$, and the results can be summed into $\mathbf{M}_{\text{tot}}$. In the end then, we are trying to compute $\min_{|\mathbf{x}|=1} \mathbf{x} \cdot \mathbf{M}_{\text{tot}}^* \mathbf{x}$, which is solved by nothing but the smallest eigenvalue of $\mathbf{M}_{\text{tot}}$ (by the Min-Max theorem). So once we know how to compute $\mathbf{M}_{\text{tot}}$, we know how to get X.

A.2.2. TEBD AND ITEBD

The time evolving block decimation (TEBD) algorithm is an algorithm that was designed for performing time evolution of a wavefunction written in MPS representation [50, 54, 96]. Roughly speaking, it is simply an application of the fact that any wavefunction can be brought into an MPS form by repeated singular value decompositions. TEBD provides an efficient way of simulating the time evolution given by the Schrödinger equation:

$$\frac{\partial}{\partial t}|\psi\rangle(t) = -iH|\psi\rangle(t),$$

where, for simplicity, we consider a time-independent Hamiltonian. The solution to this equation is formally given by

$$|\psi\rangle(t) = e^{-iHt}|\psi\rangle(0),$$

from which we define the time evolution operator $U = \exp -iHt$. Since we know how to apply operators to an MPS representation (by contraction over the physical indices), we already have the neccessary ingredients for implementing TEBD:

- Apply the time evolution operator (for an infinitesimal amount of time) to the MPS to generate a new MPS.
- This new MPS may not be in a canonical form, and may require a larger bond dimension. So we perform an SVD to truncate the resulting MPS, and repeat.

A.3. MPS FOR DENSITY MATRICES

A.3.1. OPEN SYSTEM TEBD

Loosely speaking, the MPS is representation is simply an efficient way of compressing a vector. The fact that this vector represents a wavefunction is of secondary importance. This gives us the freedom to extend the formalism and use it to compress a density matrix. To this end, we need to *vectorize* the density matrix ρ so that it becomes a vector denoted $|\rho\rangle_\sharp$, and then use the techniques described previously (mainly, SVD) to bring this vector into MPS form.

Of course, having the density matrix in MPS form is only half of the problem. The other half consists of performing time evolution on this MPS and being able to perform measurements on it. Luckily, the (i)TEBD algorithm knows how to integrate equations of the form

$$\frac{\partial}{\partial t}|v\rangle = A|v\rangle, \tag{A.24}$$

where $|v\rangle$ is a general vector. For $|v\rangle$ being a wavefunction and $|A = iH$, this is the usual time evolution of the Schrödinger equation. The time evolution of a density matrix in the presence of a Markovian environment is given by a master equation in Lindblad form:

$$\frac{\partial}{\partial t}\rho = -i\left[H,\rho\right] + \sum_\alpha \gamma_\alpha L_\alpha \rho L_\alpha^\dagger - \frac{1}{2}\left\{L_\alpha^\dagger L_\alpha, \rho\right\}, \tag{A.25}$$

where $\{\ldots\}$ denotes the anit-commutator, L_α is called a 'jump-operator' and describes the system operators that the environment couples to, and γ_α quantifies the coupling strength for jump-operator L_α. The right-hand side of this equation is encountered in different forms in the literature, making use of *superoperators*:

$$\frac{\partial}{\partial t}\rho = \underbrace{-i\left[H,\rho\right] + \underbrace{\sum_\alpha \underbrace{\gamma_\alpha L_\alpha \rho L_\alpha^\dagger - \frac{1}{2}\left\{L_\alpha^\dagger L_\alpha, \rho\right\}}_{\mathscr{D}_\alpha[\rho]}}_{\mathscr{D}[\rho]}}_{\mathscr{L}[\rho]}. \tag{A.26}$$

The superoperator $\mathscr{D}_\alpha$ is an operator acting on an operator (ρ), and is refered to as the dissipator. The operator acting on the density matrix to produce the whole right-hand side is also called the Liouvillean (superoperator), denoted $\mathscr{L}$.

Using the same vectorisation procedure that turns the operator ρ into a vector $|\rho\rangle_\sharp$, we turn a superoperator $\mathscr{L}$ into a matrix $\mathscr{L}^\sharp$. In this notation, the time evolution of the density matrix turns into

$$\frac{\partial}{\partial t}|\rho\rangle_\sharp = \mathscr{L}^\sharp|\rho\rangle_\sharp. \tag{A.27}$$

Comparing Eqns. A.24 and A.27, we can see that the time evolution of the density matrix in vectorized form is formally equivalent to the Schrödinger equation. This allows usto use the (i)TEBD algorithm for simulating mixed states.

To be more concrete about writing the density matrix as a vector in MPS format, we introduce a set of basis matrices for a local site i and denote them σ_{s_i}. If the local Hilbert space has dimension d, there are d^2 local basis matrices. For example for spin-1/2 we have $d = 2$, and the four matrices are the identity matrix plus the three Pauli matrices. In the case of a general d, we always choose the identity matrix as σ_0 and use the generators of $SU(d)$ for the remaining matrices. This choice of generators of $SU(d)$ guarantees the matrices to be traceless, which turns out to be very efficient in the evaluation of observables later. The density matrix is then written as

$$|\rho\rangle_\sharp = \sum_{s_1 \ldots s_L} \mathrm{Tr}\left(A^{[1],s_1} \cdots A^{[L],s_L}\right) |\sigma_{s_1} \ldots \sigma_{s_L}\rangle_\sharp, \tag{A.28}$$

where the 'sharping' simply has to be a consistent procedure for turning a matrix into a vector, e.g. by appending the rows or the columns. Given these basis vectors, we also have a procedure for turning any superoperator into a matrix. To do this, we first define an inner product between the sharped vectors:

$${}_\sharp\langle A|B\rangle_\sharp \equiv \frac{1}{D}\mathrm{Tr}\left(A \cdot B^\dagger\right), \tag{A.29}$$

where D is the dimension of the matrices. Notice that with our choice of basis matrices, we have ${}_\sharp\sigma_i|\sigma_j\rangle_\sharp = \delta_{ij}$. The way a superoperator is defined to act on an operator can then be carried over to the 'sharped' language, via

$$\mathscr{L}^\sharp|\rho\rangle_\sharp = |\mathscr{L}[\rho]\rangle_\sharp. \tag{A.30}$$

After all, the result of the vectorisation should be consistent with having applied the superoperator first before turning the density matrix into a vector. Numerically, to compute the matrix $\mathscr{L}^\sharp$ corresponding to the superoperator $\mathscr{L}$, we may write out its matrix elements as

$$\begin{aligned}\left(\mathscr{L}^\sharp\right)_{\alpha\beta} &= {}_\sharp\langle\sigma_\alpha|\mathscr{L}^\sharp|\sigma_\beta\rangle_\sharp \\ &= {}_\sharp\langle\sigma_\alpha|\mathscr{L}[\sigma_\beta]\rangle_\sharp \\ &= \mathrm{Tr}\left(\sigma_\alpha \cdot \left(\mathscr{L}[\sigma_\beta]\right)^\dagger\right).\end{aligned} \tag{A.31}$$

With all of these components, we are almost set for performing time evolution of mixed state density matrices. Compared to the pure states, there are minor differences in the way observables are measured on the mixed state. In the diagrammatic notation, a general measurement of an observable $\mathscr{O}$ on site number 2 of a pure state can be represented as

$\langle\psi|\mathscr{O}^{[2]}|\psi\rangle =$ 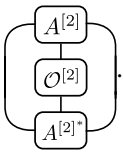

. (A.32)

For the mixed state case, let us consider the case of a spin-1/2 system. The basis matrices are chosen as discussed previously: $\sigma_{i=0,1,2,3} = \{\sigma_0, \sigma_x, \sigma_y, \sigma_z\}$. The expectation value of a local operator is written as $\langle \mathcal{O} \rangle = \mathrm{Tr}(\rho \mathcal{O})$. For the pure state case, the cyclic property of the trace led to this being equivalent to the sandwiching of the operator between two MPS copies.

In the mixed state case, the measurement of the expectation value can be captured by a superoperator $T_{\mathcal{O}}[\rho] = \mathrm{Tr}(\rho\mathcal{O})$. One may compute the matrix elements of $T^{\sharp}_{\mathcal{O}}$ using the procedure of Eq. A.31. Let us compute the operator product $\rho\mathcal{O}$ explicitly:

$$\begin{aligned} |\rho\mathcal{O}\rangle_\sharp &= T^{\sharp}_{\mathcal{O}} |\rho\rangle_\sharp \\ &= \sum_{\substack{s_1 \dots s_L \\ \alpha_1 \dots \alpha_L; \beta_1 \dots \beta_L}} \rho_{s_1 \dots s_L} \mathcal{T}^{\sharp}_{\alpha_1 \dots \alpha_L, \beta_1 \dots \beta_L} |\sigma_{\alpha_1} \dots \sigma_{\alpha_L}\rangle_\sharp \langle \sigma_{\beta_1} \dots \sigma_{\beta_L} | \sigma_{s_1} \dots \sigma_{s_L} \rangle_\sharp \\ &= \sum_{\substack{s_1 \dots s_L \\ \alpha_1 \dots \alpha_L}} \rho_{s_1 \dots s_L} \mathcal{T}^{\sharp}_{\alpha_1 \dots \alpha_L, s_1 \dots s_L} |\sigma_{\alpha_1} \dots \sigma_{\alpha_L}\rangle_\sharp . \end{aligned} \tag{A.33}$$

If we were to 'unsharp' this relation, we find

$$\rho\mathcal{O} = \sum_{\substack{s_1 \dots s_L \\ \alpha_1 \dots \alpha_L}} \rho_{s_1 \dots s_L} \mathcal{T}^{\sharp}_{\alpha_1 \dots \alpha_L, s_1 \dots s_L} \sigma_{\alpha_1} \otimes \dots \otimes \sigma_{\alpha_L}, \tag{A.34}$$

where we have restored the explicit Kronecker products between the basis matrices. Since they are traceless except for σ_0, the (unnormalized) expectation value obtained via $\mathrm{Tr}\rho\mathcal{O}$ is explicitly given by

$$\mathrm{Tr}\rho\mathcal{O} = D^L \sum_{s_1 \dots s_L} \rho_{s_1 \dots s_L} \mathcal{T}^{\sharp}_{0 \dots 0, s_1 \dots s_L}, \tag{A.35}$$

where D again is the dimension of the basis matrices, and hence $\mathrm{Tr}\sigma_0 = D$. To get the normalized expectation value, we must also compute $\mathrm{Tr}\rho$, which turns out to be equal to $D^L \rho_{0 \dots 0}$. Finally the, the normalized expecation value of $\mathcal{O}$ is given through

$$\begin{aligned} \langle \mathcal{O} \rangle &= \frac{\mathrm{Tr}\rho\mathcal{O}}{\mathrm{Tr}\rho} \\ &= \frac{1}{\rho_{0 \dots 0}} \sum_{s_1 \dots s_L} \rho_{s_1 \dots s_L} \mathcal{T}^{\sharp}_{0 \dots 0, s_1 \dots s_L}. \end{aligned} \tag{A.36}$$

A.3.2. AMPLITUDES

In the pure state case, the sandwiching of the operator between two copies of the MPS allows for full use of the canonical form. In the mixed state case, the selection of the 0-legs removes the option of sandwiching and hence is inherently an evaluation on a finitely sized system. To be able to do the contraction on a mixed state as is done for the pure state, we would have to effectively simulate the square root of the density matrix. Such a density matrix decomposition can be done in principle, via

$$\rho = \omega^\dagger \omega, \tag{A.37}$$

where ω is referred to as an amplitude for the density matrix. Notice that ω is not unique, since $U\omega$ with U a unitary matrix also produces the same density matrix. Let us specify to using a master equation with a single jump operator for simplicity:

$$\dot{\rho} = -i\left[H,\rho\right] + \gamma\left(L\rho L^\dagger - \frac{1}{2}\rho L^\dagger L - \frac{1}{2}L^\dagger L\rho\right).$$

We now substitute the decomposition of Eq. A.37 of the density matrix into amplitudes ω and ask ourselves what the equations of motion for ω look like (or equivalently, those for $\omega^\dagger$).

If we only have unitary evolution governed by the Hamiltonian H, we find

$$\dot{\omega}\omega^\dagger + \omega\dot{\omega}^\dagger = -iH\omega\,\omega^\dagger + \omega\, i\omega^\dagger H,$$

which has been written in a suggestive form, namely after a reordering of terms we have

$$(\dot{\omega} + iH\omega)\,\omega^\dagger = -\omega\left(\dot{\omega}^\dagger - i\omega^\dagger H\right).$$

Hence we are trying to solve an equation of the form $A = -A^\dagger$, which we satisfy by choosing $A = 0$, leading to $\dot{\omega} + iH\omega = 0$. This defines the time evolution for the amplitude ω, and can easily be implemented using the techniques discussed previously in this Appendix.

The problematic part is in the non-unitary evolution, governed by our jump operator L. The master equation with only the dissipative part then turns into

$$\dot{\omega}\omega^\dagger + \omega\dot{\omega}^\dagger = L\omega\omega^\dagger L^\dagger - \frac{1}{2}\omega\omega^\dagger L^\dagger L - \frac{1}{2}L^\dagger L\omega\omega^\dagger.$$

The anti-commutator between ρ and $L^\dagger L$ can be treated exaclty as the unitary part (up to a single sign change), but it is the first term of the dissipator that couples the ω and $\omega^\dagger$ terms. This prohibits an equation of motion that separates ω from $\omega^\dagger$.

A.3.3. RELATION BETWEEN ENTANGLEMENT SPECTRA

In analogy to the entanglement spectrum λ obtained from the pure state MPS, we may compute the operator space entanglement spectrum $\lambda_\sharp$ from the mixed state MPS. In this section we establish a connection between the two. For pure states, this connection turns out to be $\lambda_\sharp = \lambda \otimes \lambda$, whereas for general mixed states the connection is not as straightforward.

PURE STATES

Consider the density matrix obtained from the pure state $||\psi\rangle$, written in MPS form:

$$\rho = |\psi\rangle\langle\psi| = \sum_{\alpha\beta\gamma\delta\ldots} \mathrm{Tr}(\Gamma^{[1]}_\alpha \lambda^{[1]} \Gamma^{[2]}_\gamma \lambda^{[2]} \ldots)\, \mathrm{Tr}(\Gamma^{*[1]}_\beta \lambda^{[1]} \Gamma^{*[2]}_\delta \lambda^{[2]} \ldots) |\alpha\gamma\ldots\rangle\langle\beta\delta\ldots| \qquad \text{(A.38)}$$

where $|\alpha\rangle$ is a local basis for the spins. On the other hand, the MPS we write directly for the density marix is given as

$$\rho = \sum_{ij\ldots} \mathrm{Tr}(\Gamma_i^{\sharp[1]} \lambda^{\sharp[1]} \Gamma_j^{\sharp[2]} \lambda^{\sharp[2]} \ldots) \sigma_i \otimes \sigma_j \otimes \ldots. \tag{A.39}$$

For the purpose of relating the Γ's and λ's to their sharped counterparts, we decompose the basis for the density matrix into the local spin basis:

$$\sigma_i = \sum_{\alpha\beta} \sigma_i^{\alpha\beta} |\alpha\rangle\langle\beta|.$$

The substitution of this decomposition results in

$$\begin{aligned}
\rho &= \sum_{ij\ldots} \mathrm{Tr}(\Gamma_i^{\sharp[1]} \lambda^{\sharp[1]} \Gamma_j^{\sharp[2]} \lambda^{\sharp[2]} \ldots) \sigma_i \otimes \sigma_j \otimes \ldots \\
&= \sum_{ij\ldots} \sum_{\alpha\beta\gamma\delta\ldots} \mathrm{Tr}(\Gamma_i^{\sharp[1]} \lambda^{\sharp[1]} \Gamma_j^{\sharp[2]} \lambda^{\sharp[2]} \ldots) \sigma_i^{\alpha\beta} \sigma_j^{\gamma\delta} \ldots |\alpha\gamma\ldots\rangle\langle\beta\delta\ldots| \\
&= \sum_{\alpha\beta\gamma\delta\ldots} \mathrm{Tr}\left(\left(\sum_i \Gamma_i^{\sharp[1]} \sigma_i^{\alpha\beta} \right) \lambda^{\sharp[1]} \left(\sum_j \Gamma_j^{\sharp[2]} \sigma_j^{\alpha\beta} \right) \lambda^{\sharp[2]} \right) \ldots |\alpha\gamma\ldots\rangle\langle\beta\delta\ldots|
\end{aligned}$$

As a final step, we reexpress the pure state density matrix as

$$\begin{aligned}
\rho = |\psi\rangle\langle\psi| &= \sum_{\alpha\beta\gamma\delta\ldots} \mathrm{Tr}(\Gamma_\alpha^{[1]} \lambda^{[1]} \Gamma_\gamma^{[2]} \lambda^{[2]} \ldots) \mathrm{Tr}(\Gamma_\beta^{*[1]} \lambda^{[1]} \Gamma_\delta^{*[2]} \lambda^{[2]} \ldots) |\alpha\gamma\ldots\rangle\langle\beta\delta\ldots| \\
&= \sum_{\alpha\beta\gamma\delta\ldots} \mathrm{Tr}\Big((\Gamma_\alpha^{[1]} \otimes \Gamma_\beta^{*[1]}) (\lambda^{[1]} \otimes \lambda^{[1]}) (\Gamma_\beta^{[2]} \otimes \Gamma_\delta^{*[2]}) (\lambda^{[2]} \otimes \lambda^{[2]}) \ldots \Big) |\alpha\gamma\ldots\rangle\langle\beta\delta\ldots|
\end{aligned}$$

where we made use of the identity

$$(A \cdot B) \otimes (C \cdot D) = (A \otimes C) \cdot (B \otimes D).$$

By comparing the two expressions for the density matrix, one may identify the relation

$$\prod_n \sum_i \Gamma_i^{\sharp[n]} \sigma_i^{\alpha\beta} \cdot \lambda^{\sharp[n]} = \prod_n (\Gamma_\alpha^{[n]} \otimes \Gamma_\beta^{*[n]}) (\lambda^{[n]} \otimes \lambda^{[n]}).$$

Specifically, we therefore find

$$\lambda^{\sharp[n]} = \lambda^{[n]} \otimes \lambda^{[n]}. \tag{A.40}$$

Keep in mind that this result so far is only valid for pure states. A few important properties can be noticed here about the spectrum $\lambda^\sharp$:

- If $\lambda^{[n]}$ is fully degenerate, with even multiplicity and smallest multiplicity equal to 2, then $\lambda^{\sharp[n]}$ will also be fully degenerate with even multiplicity everywhere, but with smallest multiplicity $2 \times 2 = 4$.

- If $\lambda^{[n]}$ is fuslly non-degenerate, then $\lambda^{\sharp[n]}$ will have degenerate pairs; it will however also always have non-degenerate entries.

These statements can be summarized as: $\lambda^{\sharp[n]}$ cannot be fully degenerate unless $\lambda^{[n]}$ is, and if this is the case the smallest multiplicity in $\lambda^{\sharp[n]}$ is twice that of $\lambda^{[n]}$.

MIXED STATE

In the mixed state case, the expression for the density matrix in terms of $\Gamma^{\sharp}$ and $\lambda^{\sharp}$ is still valid. However, the mixed state expression for a general density matrix is

$$\rho = \sum_m c_m |\psi_m\rangle\langle\psi_m|$$

Let us give the Γ and λ matrices yet another index, with the notation $\Gamma^{[n,m]}$ to indicate that this is the Γ matrix on site n belonging to the MPS of the wavefunction $|\psi_m\rangle$. In this notation, we expand

$$\rho = \sum_m \sum_{\alpha\beta\ldots} \mathrm{Tr}\left(c_m \left(\Gamma_\alpha^{[1,m]} \otimes \Gamma_\beta^{*[1,m]}\right)(\lambda^{[1,m]} \otimes \lambda^{[1,m]})\ldots\right) |\alpha\ldots\rangle\langle\beta\ldots|$$

Comparing this to the expression for the sharped density matrix in the previous section, we find

$$\prod_n \sum_i \Gamma_i^{\sharp[n]} \sigma_i^{\alpha\beta} \cdot \lambda^{\sharp[n]} = \sum_m c_m \prod_n (\Gamma_\alpha^{[n,m]} \otimes \Gamma_\beta^{*[n,m]})(\lambda^{[n,m]} \otimes \lambda^{[n,m]}).$$

The relation between $\lambda^{\sharp}$ and the $\lambda^{[n,m]}$ is not straightforwardly extracted from this expression.

B

REPRESENTATION THEORY

We review the basic concepts of linear and projective representation theory, and discuss the consequences of projective representations for the classification of symmetry protected topological states.

B.1. REPRESENTATION THEORY

In this section we first review basic concepts of group theory and representation theory. These are the prerequisites for discussing projective representation in section B.2. We finish this appendix by a short recap of the projective symmetries applied to matrix product states, demonstrating how they allow for classifying topological phases in 1D. This last section can also be read on its own, without making explicit use of the rigorous definitions in the following sections. Some mathematical rigour has been sacrificed in this chapter in favor of highlighting the logical connections between the concepts, without introducing too many extra bells and whistles.

B.1.1. LINEAR REPRESENTATION THEORY

Let us consider a collection of objects g_i that we collect into another object $\mathcal{G} = \{g_i\}$. In addition, we define some operation by which we can combine the objects g_i. Let us denote this operation by $\circ$, s.t. we may write $\tilde{g} = g_i \circ g_j$ for g_i and g_j both in $\mathcal{G}$. Then $\mathcal{G}$ together with the operator $\circ$ define a *group* if they satisfies the following four group axioms:

1. Closure: If g_i and g_j are in $\mathcal{G}$, then so is $g_i \circ g_j$.
2. Associativity: $g_i \circ (g_j \circ g_k) = (g_i \circ g_j) \circ g_k$ for g_i, g_j, g_k all in $\mathcal{G}$
3. Identity element: there is an element e in $\mathcal{G}$, s.t. $e \circ g_i = g_i \circ e = g_i$.
4. Inverse element: for every element g_i in $\mathcal{G}$, there is an element g_i^{-1} s.t. $g_i \circ g_i^{-1} = g_i^{-1} \circ g_i = e$.

The number of elements g_i in the group is refered to as the *order* of the group, denoted $|\mathcal{G}|$, and the group is called finite if this number is finite. Additionaly, a group is called Abelian if for all its elements we have $g_i \circ g_j = g_j \circ g_i$.

A linear representation is now defined by assigning a matrix $R(g_i)$ (with non-zero determinant) to each element g_i of a group. This assignment comes in the form of a homomorphism:

$$g_i \in \mathcal{G} \mapsto R(g_i) \in GL(n). \tag{B.1}$$

The space $GL(n)$ represents the space of invertible linear operators on an n-dimensional vector space, or in other words the space of $n \times n$ matrices with non-zero determinant. This homomorphism defines a valid linear representation if it satisfies the following three properties:

1. The identity element $e \in \mathcal{G}$ is mapped to the identity matrix $R(e) = \hat{1}$.

2. The inverse of an element is mapped to the matrix inverse: $R(g^{-1}) = \big(R(g)\big)^{-1}$.
3. The group action translates to matrix multiplication: $R(g_1 \circ g_2) = R(g_1)R(g_2)$.

We call n the dimension of the representation. It is important to state clearly here that a representation is not unique. It is possible to find different sets of $n \times n$ matrices for the group elements g that all satisfy the requirements of being a representation. We therefore need a way of determining which representations are equivalent. This equivalence relation is given by the usual matrix similarity, i.e. two representations $\tilde{R}$ and R are equivalent if there exists a matrix $A \in GL(n)$ such that

$$\tilde{R}(g) = AR(g)A^{-1} \tag{B.2}$$

for all elements $g \in \mathcal{G}$ simultaneously, i.e. A does not depend on g. Notice that for finite groups there is an easy way of checking whether two representations are equivalent. Namely, the similarity transformation does not affect the trace of the matrices (also called the characters of a representation), and so if two representations have the same characters they are equivalent.

This brings us to the concept of reducability of a representation. A representation R is said to be **irreducible** if is *not* equivalent to another representation $\tilde{R}$ that is block-diagonal: $\tilde{R} = \tilde{R}_1 \oplus \tilde{R}_2$, where $\tilde{R}_1$ and $\tilde{R}_2$ themselves are also irreducible. By extension, a representation is **reducible** if it is expressible as a direct sum of irreducible representations. In other words, a representation is reducible if it is equivalent to a representation that is block-diagonal, and irreducible if it is not.

It is in general not an easy problem to find all the irreducible representations of a symmetry group, and it is easiest to simply look them up. There are some simpler cases in which one can deduce the irreducible representations directly, as is the case for finite Abelian groups: all irreducible representations of such groups are 1 dimensional. This is a direct consequence of Schur's lemmas [145]. When combined with the knowledge that $\sum_i n_i^2 = |\mathcal{G}|$, with n_i the dimension of irreducible representation R_i, means that a finite Abelian group of order m has exactly m irreducible representations of dimension one. An example of this is given in Table B.1, for the symmetry group $\mathcal{G} = \mathbb{Z}_2 \times \mathbb{Z}_2 = \{e, g_x, g_z, g_{xz}\}$.

	$\hat{1}$	g_x	g_z	g_{xz}
R_1	1	1	1	1
R_2	1	1	-1	-1
R_3	1	-1	1	-1
R_4	1	-1	-1	1

Table B.1: All the irreducible representations R_i of the group $\mathbb{Z}_2 \times \mathbb{Z}_2$. Since this group is finite and Abelian all its representations have dimension $n_i = 1$.. In addition, since it has order 4, we know from $\sum_i n_i^2 = 4$ that there are 4 irreducible representations. It is easily confirmed that the choices in this table indeed satisfy the requirements for them to be valid linear representations.

B.1.2. PROJECTIVE REPRESENTATION THEORY

Let us turn our attention from the linear representations to the more general projective representations. The material in this section is based on Ref. [146]. We need to first introduce the concepts of quotient groups and the center of a group.

Given a group $\mathcal{G}$ and a subgroup[1] $\mathcal{H} \subset \mathcal{G}$, the quotient group $\mathcal{G}/\mathcal{H}$ is defined by all the elements $\{gH\}$ for $g \in \mathcal{G}$, with $gH = \{g \circ h\}$ for all $h \in \mathcal{H}$. This can be examplified by taking as the group $\mathcal{G} = \mathbb{Z}$, the group of all the integers, and as a subgroup $\mathcal{H} = 2\mathbb{Z}$, the subgroup of all the even integers. The group operation $g \circ h$ is simply given by addition of the two numbers. Then the quotient group $\mathcal{G}/\mathcal{H} = \mathbb{Z}/2\mathbb{Z}$ consists of two different elements, namely all the even integers and all the odd integers. This is therefore equivalent to $\mathbb{Z}_2 = \{0, 1\}$, the cyclic group of order 2.

The center of a group $\mathcal{G}$ is denoted $Z(\mathcal{G})$, and is defined by the collection $\{z \in \mathcal{G}\}$ s.t. $z \circ g = g \circ z$ for all g in $\mathcal{G}$. In words, these are all the elements in $\mathcal{G}$ that commute with the whole of $\mathcal{G}$. For an Abelian group, this means that the center is simply the entire group itself.

We can now introduce projective representations by analogy to the linear ones in the previous section. Again, we associate to each element g of a group $\mathcal{G}$ a matrix denoted $\rho(g)$. Instead of $\rho(g)$ beign a member of $GL(n)$, it is now a member of the projective linear group $PGL(n)$, defined as $PGL = GL/Z(GL)$. The representation ρ is a correct projective representation if it satisfies the following three conditions:

1. The identity element $e \in \mathcal{G}$ is mapped to the identity matrix $\rho(e) = \hat{1}$.
2. The inverse of an element is mapped to the matrix inverse: $\rho(g^{-1}) = \big(\rho(g)\big)^{-1}$.
3. The group action translates to matrix multiplication, *up to a factor* $c(g_1, g_2)$: $\rho(g_1 \circ g_2) = c(g_1, g_2)\rho(g_1)\rho(g_2)$.

The difference to the linear case is the factor $c(g_1, g_2)$ in the third condition. This factor is an element of the center of the general linear group $Z(GL)$, which is equivalent to the group $U(1)$ and hence is a phase factor, $c(g_1, g_2) = e^{i\phi(g_1, g_2)}$. These phases are also referred to as a Schur multipliers, or more commonly as the factor set on $\mathcal{G}$. Obviously, if $\phi(g_1, g_2) = 0$ for all g_1, g_2 in $\mathcal{G}$, the projective representation reduced to a linear one.

The equivalence of projective representations carries over from the linear representations too, again up to a phase factor $\alpha(g) \in Z(GL)$. As such, an n-dimensional representation $\rho_2(g)$ is equivalent to $\rho_1(g)$ if there is a matrix A from $GL(n)$ for which

$$\rho_2(g) = \alpha(g) A \rho_1(g) A^{-1}. \tag{B.3}$$

[1] A subgroup of a group $\mathcal{G}$ is simply a subset of the elements in $\mathcal{G}$ that themselves form a group.

Also by analogy, a projective representation ρ is reducible if it is equivalent to another projective representation $\tilde{\rho}$ that is in block diagonal form, i.e.

$$\tilde{\rho} = \begin{pmatrix} \tilde{\rho}_1 & 0 \\ 0 & \tilde{\rho}_2 \end{pmatrix}.$$

The factor set (phases $\phi(g_1, g_2)$) has three important properties. The first is a direct consequence of the fact that $PGL(n)$ is a group, and hence its elements are associative. This means that $\rho(g_1)\big(\rho(g_2)\rho(g_3)\big) = \big(\rho(g_1)\rho(g_2)\big)\rho(g_3)$ and hence

$$e^{i\phi(g_1,g_2g_3)}e^{i\phi(g_2,g_3)} = e^{i\phi(g_1,g_2)}e^{i\phi(g_1g_2,g_3)}, \tag{B.4}$$

an equation known as the 2-cocycle equation. Another property is the 2-coboundary, stemming from the phase factor $\alpha(g)$ in Eq. B.3. This means that we can choose a new representation $\tilde{\rho}(g) = \alpha(g)\rho(g)$, and obtain a new factor set as

$$e^{i\tilde{\phi}(g_1,g_2)} = \frac{\alpha(g_1g_2)}{\alpha(g_1)\alpha(g_2)}e^{i\phi(g_1,g_2)}. \tag{B.5}$$

The factor involving the ratio of the α's is known as the 2-coboundary, and is used to define an equivalence class between factor sets (i.e. two factor sets are equivalent if there exists a 2-coboundary relation them as in Eq. B.5).

Lastly, the equivalence classes of factor sets turn out to form a group themselves. This group is Abelian and is named the second cohomology group $H^2(\mathcal{G}, U(1))$ over $U(1)$, in which the group operation corresponds to the multiplication of the phase factors. This group also has an identity element (the 'trivial class'), which is given by all the phase factors that can be transformed to $\phi(g_1, g_2) = 0$ using an appropriate 2-coboundary [49, 145, 146].

For our running example of the $\mathbb{Z}_2 \times \mathbb{Z}_2$, we find the 2-dimensional irreducible representations

$$\rho(e) = \hat{1}, \quad \rho(g_x) = \sigma_x, \quad \rho(g_z) = \sigma_z, \quad \rho(g_zx) = i\sigma_y \tag{B.6}$$

where the σ matrices are the 2×2 Pauli matrices. The second cohomology group of $\mathbb{Z}_2 \times \mathbb{Z}_2$ is $H^2(\mathbb{Z}_2 \times \mathbb{Z}_2, U(1)) = \mathbb{Z}_2$, meaning that there are only two equivalence classes of the factor set: the trivial one and a non-trivial one.

B.2. PROJECTIVE REPRESENTATIONS ON MPS

This section serves as a pragmatic summary of the projective representations discussed in the first part of this appendix. It can either be read independently, or may serve as an example and summary of the concepts discussed before.

We base our discussion of projective symmetries on a wavefunction represented in an MPS form, and follow the general arguments as outlined in Ref. [147]. We consider a

system that has a symmetry group $\mathcal{G}$, and a groundstate that is invariant under the elements $g \in \mathcal{G}$. We assume that the symmetry operations are local, and that local states (denoted as $|s\rangle$) transform as

$$|s'\rangle \to \sum_s \left(R_g\right)_{s's} |s\rangle, \tag{B.7}$$

with $\left(R_g\right)$ the unitary matrix representing the symmetry operation g. If the MPS representation of the wavefunction uses matrices denoted M, the action of the symmetry operation be represented by a change of the MPS matrices M as:

$$\tilde{M}^{s'} = \sum_s \left(R_g\right)_{s's} M^s = e^{i\phi_g} U_g^{-1} M^{s'} U_g. \tag{B.8}$$

The matrix U_g is the representation of the symmetry on the level of the MPS, and is unique up to a $U(1)$ phase as long as the MPS is in canonical form.

Now let us consider the successive action of two symmetry elements g and h:

$$\tilde{M}^{s''} = \sum_{s,s'} \left(R_g\right)_{s''s'} \left(R_h\right)_{s's} M^s = e^{i\phi_g} e^{i\phi_h} U_h^{-1} U_g^{-1} M^{s''} U_g U_h. \tag{B.9}$$

Instead, we could also have considered the action of the symmetry element gh directly, which would have led to:

$$\tilde{M}^{s''} = \sum_s \left(R_{gh}\right)_{s''s} M^s = e^{i\phi_{gh}} U_{gh}^{-1} M^{s''} U_{gh}. \tag{B.10}$$

By comparing these last two expression, we must have $\phi_{gh} = \phi_g + \phi_h$ and

$$U_g U_h = e^{i\phi(g,h)} U_{gh}, \tag{B.11}$$

where $\phi(g,h)$ is a general $U(1)$ phase, showing that the U's form a projective representation of the symmetry group $\mathcal{G}$. This phase factor contains the information on the topological nature of the wavefunction. We discuss two of its general properties.

First, let us make use of the associativity of the group elements and calculate the product of three elements g_1, g_2 and g_3 all in $\mathcal{G}$:

$$U_{g_1} U_{g_2} U_{g_3} = e^{i\phi(g1,g2)} U_{g_1 g_2} U_{g_3} = e^{i\left(\phi(g_1,g_2)+\phi(g_1 g_2,g_3)\right)} U_{g_1 g_2 g_3},$$

versus

$$U_{g_1} U_{g_2} U_{g_3} = e^{i\phi(g2,g3)} U_{g_1} U_{g_2 g_3} = e^{i\left(\phi(g_2,g_3)+\phi(g_1,g_2 g_3)\right)} U_{g_1 g_2 g_3}.$$

Comparing these results we find that the phase factors satisfy the *2-cocycle* condition:

$$\phi(g1, g_2 g_3) - \phi(g_1 g_2, g_3) + \phi(g_2, g_3) - \phi(g_1, g_2) = 0. \tag{B.12}$$

Second, there is a degree of freedom in Eq. B.8 that allows us to multiply the matrix U_g by an arbitrary phase factor $e^{i\beta(g)}$. This changes the phase factor in Eq. B.11 according to

$$\phi(g,h) \to \phi(g,h) + \left(\beta(h) - \beta(gh) + \beta(g)\right),$$

where the term in parenthesis is called the *2-coboundary*. The new phase factor $\phi(g,h)$ is equivalent to the old phase factor up to 2-coboundaries, and both describe the same topological phase.

More rigorously, the set of phase functions satifsying the 2-cocyle condition can be divided into equivalence classes under the equivalence relation that is the 2-coboundary. These equivalence classes themselves form a group, that is known as the second cohomology group of the group $\mathcal{G}$ over $U(1)$, denoted $H^2(\mathcal{G},U(1))$. Hence, given a symmetry group $\mathcal{G}$, its second cohomology tells us exactly what the allowed topological phases are characterized by. The spin-1 Heisenberg model discussed in Chapter 2 is a good example. The symmetry group we discussed was $\mathbb{Z}_2\times\mathbb{Z}_2$ whose second cohomology group $H^2(\mathbb{Z}_2\times\mathbb{Z}_2,U(1))=\mathbb{Z}_2$, exactly indicating that there are *two* distinct phases preserving this symmetry: the Haldane phase and the Large-D phase. The Large-D phase corresponds to the trivial choice of phase functions, i.e. the equivalence class to which the choice $\phi(g,h)=0$ for all g,h belongs. In general for the 1-dimensional symmetry protected topological phases, it was shown that the second cohomology group is able to fully classify the different allowed topological phases [49].

REFERENCES

[1] P. W. Anderson, Science **177**, 393 (1972), URL.

[2] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics* (Cambridge University Press, 2000).

[3] J. M. Yeomans, *Statistical Mechanics of Phase Transitions* (Oxford University Press, 1992).

[4] C. Kittel, *Introduction to Solid State Physics* (Wiley, 2004), 8th ed.

[5] C. J. Foot, *Atomic Physics* (Oxford University Press, 2005).

[6] P. W. Anderson, *Basic Notions of Condensed Matter Physics* (Benjamin/Cummings, 1984).

[7] V. L. Ginzburg and L. Landau, Zh. Eksp. Teor. Fiz. 20 **1064** (1950).

[8] J. Zaanen, *The Classical Condensates – from crystals to Fermi-liquids* (University of Leiden, 1996).

[9] M. Z. Hasan and C. L. Kane, Rev. Mod. Phys. **82**, 3045 (2010), URL.

[10] X.-L. Qi and S.-C. Zhang, Rev. Mod. Phys. **83**, 1057 (2011), URL.

[11] T. Senthil, Annual Review of Condensed Matter Physics **6**, 299 (2015), URL.

[12] N. Laflorencie, arXiv e-prints p. 1512.03388 (2015), URL.

[13] M. Rigol, V. Dunjko, and M. Olshanii, Nature **452**, 854 (2008), URL.

[14] Y. Bahri, R. Vosk, E. Altman, and A. Vishwanath, Nat Commun **6** (2015), URL.

[15] H. Goldstein, *Classical Mechanics* (Addison-Wesley, 1980), 2nd ed.

[16] P. W. Anderson, Phys. Rev. **109**, 1492 (1958), URL.

[17] R. Nandkishore and D. A. Huse, Annual Review of Condensed Matter Physics **6**, 15 (2015), URL.

[18] V. Oganesyan and D. A. Huse, Physical Review B **75**, 155111 (2007), URL.

[19] M. Serbyn, Z. Papić, and D. A. Abanin, Phys. Rev. Lett. **111**, 127201 (2013), URL.

[20] G. Lindblad, Comm. Math. Phys **48**, 119 (1976), URL.

[21] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2001), 1st ed.

[22] U. Weiss, *Quantum Dissipative Systems*, vol. 2 of *Series in Modern Condensed Matter Physics* (World Scientific, 1999).

[23] R. S. M. J. G. Carbonell and T. M. Mitchell, *An overview of machine learning*, no. 1 in Machine Learning: An Artificial Intelligence Approach (Springer, 1983).

[24] S. B. Kotsiantis, *Supervised machine learning: A review of classification techniques*, Frontiers in Artificial Intelligence and Applications (IOS Press, 2007).

[25] I. Affleck, J. Phys.: Condens. Matter. **1**, 3047 (1989), URL.

[26] R. Kubo, Phys. Rev. **87**, 568 (1952), URL.

[27] P. W. Anderson, Phys. Rev. **86**, 694 (1952), URL.

[28] H. Bethe, Zeitschrift für Physik **71**, 205 (1931), URL.

[29] J. M. Kosterlitz and D. J. Thouless, Journal of Physics C: Solid State Physics **6**, 1181 (1973), URL.

[30] F. D. M. Haldane, Physics Letters A **93**, 464 (1983), URL.

[31] F. D. M. Haldane, Phys. Rev. Lett. **50**, 1153 (1983), URL.

[32] R. Botet and R. Jullien, Phys. Rev. B **27**, 613 (1983), URL.

[33] J. B. Parkinson and J. C. Bonner, Phys. Rev. B **32**, 4703 (1985), URL.

[34] M. P. Nightingale and H. W. J. Blöte, Phys. Rev. B **33**, 659 (1986), URL.

[35] E. Lieb, T. Schultz, and D. Mattis, Annals of Physics **16**, 407 (1961).

[36] I. Affleck and E. H. Lieb, Letters in Mathematical Physics **12**, 57 (1986), URL.

[37] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Physical Review Letters **59**, 799 (1987), URL.

[38] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Communications in Mathematical Physics **115**, 477 (1988), URL.

[39] C. K. Majumdar and D. K. Ghosh, Journal of Mathematical Physics **10** (1969).

[40] N. Schuch, D. Pérez-García, and I. Cirac, Phys. Rev. B **84**, 165139 (2011), URL.

[41] Z.-C. Gu and X.-G. Wen, Physical Review B **80**, 155131 (2009), URL.

[42] F. Pollmann and A. Turner, Phys. Rev. B **86**, 125441 (2012), URL.

[43] K. Wierschem and P. Sengupta, Modern Physics Letters B **28**, 1430017 (2014), URL.

[44] M. den Nijs and K. Rommelse, Phys. Rev. B **40**, 4709 (1989), URL.

[45] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Bañuls, L. Pollet, I. Bloch, et al., Science **334**, 200 (2011), URL.

[46] D. Pérez-García, M. M. Wolf, M. Sanz, F. Verstraete, and J. I. Cirac, Phys. Rev. Lett. **100**, 167202 (2008), URL.

[47] F. Pollmann, A. M. Turner, E. Berg, and M. Oshikawa, Phys. Rev. B **81**, 064439 (2010), URL.

[48] A. M. Turner, F. Pollmann, and E. Berg, Phys. Rev. B **83**, 075102 (2011), URL.

[49] X. Chen, Z.-C. Gu, and X.-G. Wen, Phys. Rev. B **83**, 035107 (2011), URL.

[50] G. Vidal, Phys. Rev. Lett. **98**, 070201 (2007), URL.

[51] M. B. Hastings, Journal of Statistical Mechanics: Theory and Experiment **2007**, P08024 (2007), URL.

[52] D. Gottesman and M. B. Hastings, New Journal of Physics **12**, 025002 (2010), URL.

[53] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Phys. Rev. Lett. **100**, 030504 (2008), URL.

[54] G. Vidal, Phys. Rev. Lett. **91**, 147902 (2003), URL.

[55] F. Pollmann, E. Berg, A. M. Turner, and M. Oshikawa, Phys. Rev. B **85**, 075125 (2012), URL.

[56] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Rev. Mod. Phys. **81**, 865 (2009), URL.

[57] H. Li and F. D. M. Haldane, Phys. Rev. Lett. **101**, 010504 (2008), URL.

[58] T. H. Hsieh and L. Fu, Phys. Rev. Lett. **113**, 106801 (2014), URL.

[59] R. Thomale, D. P. Arovas, and B. A. Bernevig, Phys. Rev. Lett. **105**, 116805 (2010), URL.

[60] R. Lundgren, V. Chua, and G. A. Fiete, Phys. Rev. B **86**, 224422 (2012), URL.

[61] S. Vijay and L. Fu, Phys. Rev. B **91**, 220101 (2015), URL.

[62] B. N. M. Fannes and R. Werner, Commun. Math. Phys **144**, 443 (1992), URL.

[63] A. Chandran, V. Khemani, and S. L. Sondhi, Phys. Rev. Lett. **113**, 060501 (2014), URL.

[64] E. M. Stoudenmire and D. J. Schwab, arXiv e-prints p. 1605.05775 (2016), URL.

[65] G. Carleo and M. Troyer, arXiv e-prints p. 1606.02318 (2016), URL.

[66] L. Wang, arXiv e-prints p. 1606.00318 (2016), URL.

[67] J. Carrasquilla and R. G. Melko, arXiv e-prints **abs/1605.01735** (2016).

[68] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Under review in Nature Physics (2016).

[69] F. Rosenblatt, Report 85-460-1, Cornell Aeronautical Laboratory (1957).

[70] S. O. Haykin, *Neural Networks: A Comprehensive Foundation* (Prentice Hall, 1998).

[71] X. L. Qi, H. Katsura, and A. W. W. Ludwig, Phys. Rev. Lett. **108**, 1 (2012), URL.

[72] P. Calabrese and A. Lefevre, Phys. Rev. A **78**, 032329 (2008), URL.

[73] V. Alba, M. Haque, and A. M. Läuchli, Phys. Rev. Lett. **108**, 227201 (2012), URL.

[74] Z.-C. Yang, C. Chamon, A. Hamma, and E. R. Mucciolo, Phys. Rev. Lett. **115**, 267206 (2015), URL.

[75] S. D. Geraedts, R. Nandkishore, and N. Regnault, Phys. Rev. B **93**, 174202 (2016), URL.

[76] H. Pichler, G. Zhu, A. Seif, P. Zoller, and M. Hafezi, arXiv e-prints p. 1605.08624 (2016), URL.

[77] R. Thomale, D. P. Arovas, and B. A. Bernevig, Phys. Rev. Lett. **105**, 116805 (2010), URL.

[78] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, Rev. Mod. Phys. **80**, 517 (2008), URL.

[79] N. Laflorencie, arXiv e-prints p. 1512.03388 (2015), URL.

[80] F. K. Pearson, Philosophical Magazine Series 6 **2**, 559 (1901).

[81] M. Nielsen, *Neural Networks and Deep Learning* (Determination Press, 2015).

[82] E. P. L. van Nieuwenburg and S. D. Huber, Phys. Rev. B **90**, 075141 (2014), URL.

[83] X.-G. Wen, Advances in Physics **44**, 405 (1995), URL.

[84] S. Yan, D. A. Huse, and S. R. White, Science **332**, 1173 (2011), URL.

[85] C. Castelnovo and C. Chamon, Phys. Rev. B **76**, 184442 (2007), URL.

[86] M. B. Hastings, Phys. Rev. Lett. **107**, 210501 (2011), URL.

[87] O. Viyuela, A. Rivas, and M. A. Martin-Delgado, Phys. Rev. B **86**, 155140 (2012), URL.

[88] A. Kitaev, AIP Conference Proceedings pp. 22–30 (2009).

[89] S. Ryu, A. P. Schnyder, A. Furusaki, and A. W. W. Ludwig, New Journal of Physics **12**, 065010 (2010).

[90] X. Chen, Z.-X. Liu, and X.-G. Wen, Phys. Rev. B **84**, 235141 (2011), URL.

[91] X. Chen, Z.-C. Gu, and X.-G. Wen, Phys. Rev. B **84**, 235128 (2011), URL.

[92] A. Rivas, O. Viyuela, and M. A. Martin-Delgado, Phys. Rev. B **88**, 155141 (2013), URL.

[93] J. E. Avron, M. Fraas, G. M. Graf, and O. Kenneth, New Journal of Physics **13**, 053042 (2011), URL.

[94] J. E. Avron, M. Fraas, and G. M. Graf, Journal of Statistical Physics **148**, 800 (2012), URL.

[95] M. Zwolak and G. Vidal, Phys. Rev. Lett. **93**, 207205 (2004), URL.

[96] R. Orús and G. Vidal, Phys. Rev. B **78**, 155117 (2008), URL.

[97] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Phys. Rev. Lett. **93**, 207204 (2004), URL.

[98] K. Kim, M.-S. Chaing, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe, Nature pp. 590–593 (2010), URL.

[99] D. M. Basko, L. I. Aleiner, and B. L. Altshuler, Problems of Condensed Matter Physics pp. 50–70 (2006), URL.

[100] A. Pal and D. A. Huse, Phys. Rev. B **82**, 174411 (2010), URL.

[101] R. Vosk and E. Altman, Phys. Rev. Lett. **110**, 067204 (2013), URL.

[102] M. Serbyn, Z. Papić, Z., and D. A. Abanin, Physical Review Letters **110**, 1 (2013), URL.

[103] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, Phys. Rev. Lett. **113**, 107204 (2014), URL.

[104] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, Science **349**, 842 (2015), URL.

[105] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe (2015), URL.

[106] J. Z. Imbrie, Journal of Statistical Physics **163**, 998 (2016), URL.

[107] W. De Roeck, F. Huveneers, M. Müller, and M. Schiulaz, Phys. Rev. B **93**, 014203 (2016), URL.

[108] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan, Phys. Rev. Lett. **42**, 673 (1979), URL.

[109] V. Ros, M. Müller, and A. Scardicchio, Nuclear Physics B **891**, 420 (2015), URL.

[110] D. A. Huse, R. Nandkishore, and V. Oganesyan, Phys. Rev. B **90**, 174202 (2014), URL.

[111] L. Rademaker and M. Ortuño, Phys. Rev. Lett. **116**, 010404 (2016), URL.

[112] A. Chandran, I. H. Kim, G. Vidal, and D. A. Abanin, Phys. Rev. B **91**, 085425 (2015), URL.

B

[113] H. Kim and D. A. Huse, Phys. Rev. Lett. **111**, 127205 (2013), URL.

[114] E. H. Lieb and D. W. Robinson, Comm. Math. Phys (1972).

[115] J. H. Bardarson, F. Pollmann, and J. E. Moore, Phys. Rev. Lett. **109**, 017202 (2012), URL.

[116] F. Giraldi and F. Petruccione, Physical Review A **88**, 042102 (2013), URL.

[117] A. G. Kofman and G. Kurizki, Phys. Rev. Lett. **93**, 130406 (2004), URL.

[118] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2002).

[119] S. Johri, R. Nandkishore, and R. Bhatt, Physical Review Letters **114**, 117401 (2015), URL.

[120] M. H. Fischer, M. Maksymenko, and E. Altman, Phys. Rev. Lett. **116**, 160401 (2016), URL.

[121] R. Nandkishore, S. Gopalakrishnan, and D. A. Huse, Phys. Rev. B **90**, 064203 (2014), URL.

[122] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, Science **352**, 1547 (2016), URL.

[123] P. Bordia, H. Lüschen, U. Schneider, M. Knap, and I. Bloch, arXiv p. 1607.07868 (2016), URL.

[124] F. M. Cucchietti, J. P. Paz, and W. H. Zurek, Physical Review A **72**, 052113 (2005), URL.

[125] J. Lages, V. V. Dobrovitski, M. I. Katsnelson, H. A. De Raedt, and B. N. Harmon, Phys. Rev. E **72**, 26225 (2005), URL.

[126] S. Camalet and R. Chitra, Physical Review B **75**, 094434 (2007), URL.

[127] D. Rossini, T. Calarco, V. Giovannetti, S. Montangero, and R. Fazio, Physical Review A **75**, 032333 (2007), URL.

[128] Y.-C. Ou and H. Fan, Journal of Physics A: Mathematical and Theoretical **40**, 2455 (2007), URL.

[129] N. Wu, A. Nanduri, and H. Rabitz, Physical Review A **89**, 062105 (2014), URL.

[130] N. Wu, N. Fröhling, X. Xing, J. Hackmann, A. Nanduri, F. B. Anders, and H. Rabitz, Physical Review B **93**, 035430 (2016), URL.

[131] J. Restrepo, S. Camalet, and R. Chitra, EPL (Europhysics Letters) **101**, 50005 (2013).

[132] E. A. Winograd, M. J. Rozenberg, and R. Chitra, Phys. Rev. B **80**, 214429 (2009), URL.

[133] H. Brox, J. Bergli, and Y. M. Galperin, Physical Review A - Atomic, Molecular, and Optical Physics **85**, 1 (2012), URL.

[134] M. Mehta, *Random Matrices*, 9780120884094 (Academic Press, 2004), 3rd ed.

[135] M. Serbyn and J. E. Moore, arXiv p. 1508.07293 (2015), URL.

[136] R. Fossion, G. T. Vargas, V. Velazquez, and J. C. L. Vieyra, Journal of Physics: Conference Series **578** (2013).

[137] Y. Y. Atas, E. Bogomolny, O. Giraud, and G. Roux, Phys. Rev. Lett. **110**, 084101 (2013), URL.

[138] M. Biondi, E. P. L. van Nieuwenburg, G. Blatter, S. D. Huber, and S. Schmidt, Phys. Rev. Lett. **115**, 143601 (2015), URL.

[139] E. P. L. van Nieuwenburg, S. D. Huber, and R. Chitra, arXiv e-prints (2016).

[140] R. Orús, Annals of Physics **349**, 117 (2014).

[141] G. M. Crosswhite and D. Bacon (2008).

[142] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010), URL.

[143] M. B. Plenio, Phys. Rev. Lett. **95**, 090503 (2005), URL.

[144] U. Schollwöck, Annals of Physics **326**, 96 (2011), URL.

[145] I. Schur, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin pp. 406–432 (1905).

[146] X. Chen, Z.-C. Gu, Z.-X. Liu, and X.-G. Wen, Phys. Rev. B **87**, 155114 (2013), URL.

[147] T. Morimoto, H. Ueda, T. Momoi, and A. Furusaki, Phys. Rev. B **90**, 235111 (2014), URL.

LIST OF PUBLICATIONS

1. J. Gudat, C. Bonato, **E.P.L. van Nieuwenburg**, S. Thon, H. Kim, P.M. Petroff, M.P. van Exter and D. Bouwmeester *Permanent tuning of quantum dot transitions to degenerate microcavity resonance*, Applied Physics Letters **98**, 121111 (2011).

2. C. Bonato, **E.P.L. van Nieuwenburg**, J. Gudat, S. Thon, H. Kim, M.P. van Exter and D. Bouwmeester *Strain-tuning of quantum dot optical transitions via laser-induced surface defects*, Physical Review B **84**, 075306 (2011).

3. **E.P.L. van Nieuwenburg**, J.M. Edge, J.P. Dahlhaus, J. Tworzydlo, C.W.J. Beenakker *Metal–topological-insulator transition in the quantum kicked rotator with Z_2 symmetry*, Physical Review B **85**, 165131 (2012).

4. M. Tovmasyan, **E.P.L. van Nieuwenburg**, S.D. Huber *Geometry-induced pair condensation*, Physical Review B rapid **88**, 220510 (2013).

5. **E.P.L. van Nieuwenburg**, S.D. Huber, *Classification of mixed state topology in 1D*, Physical Review B **90**, 075141 (2014).

6. M. Biondi, **E.P.L. van Nieuwenburg**, S.D. Huber, G. Blatter, S. Schmidt *Polaritonic Flatbands*, Physical Review Letter **115**, 143601 (2015).

7. **E.P.L. van Nieuwenburg**, S.D. Huber, R. Chitra, *Single spin probe of Many-Body Localization*, ArXiv e-Print, **submitted to PRB Rapid Communication**.

8. **E.P.L. van Nieuwenburg**, Y-H. Liu, S.D. Huber, *Learning phase transitions by confusion*, **under review for Nature Physics**.

9. **E.P.L. van Nieuwenburg**, S.D. Huber, *Thermalization of a Many-Body Localized system coupled to baths*, **in preparation**.

10. Paolo Molignini, **E.P.L. van Nieuwenburg**, R. Chitra, *Floquet Majoranas*, **in preparation**.

Curriculum Vitae

Personal data

Name	Everard Pieter Lodewijk van Nieuwenburg
Date of birth	10th of March 1990 (Rotterdam, Netherlands)
Nationality	Netherlands

Education

2012 - 2016	Doctoral studies at ETH Zürich under supervision of Prof. Dr. Sebastian Huber
2010 - 2012	MSc. studies at Leiden University, Netherlands
	MSc. thesis under supervision of Prof. Dr. Carlo Beenakker
	Title: *Metal–Topological-Insulator transition in a quantum kicked rotator with $\mathbb{Z}_2$ symmetry*
2007 - 2010	BSc. studies at Leiden University, Netherlands
	BSc. thesis under supervision of Prof. Dr. Dirk Bouwmeester
	Title: *Solid state CQED: Quantum dots in microcavities*
2007 - 2012	Undergraduate studies at Leiden University, Netherlands
2001 - 2007	Secondary school at Stedelijk Gymnasium in Schiedam, Netherlands

ACKNOWLEDGEMENTS

Without the presence and involvement of all the people I have interacted with over the last few years, this thesis would not have been possible. First and foremost, I would like to express my deepest gratitude to my supervisor Prof. Dr. Sebastian Huber, not in the least for taking a chance on me as his first PhD student. It has been an absolute pleasure working with him, being given the freedom to pursue many interesting topics as well as learning from his expertise on Physics and his viewpoints on the way the world works.

Furthermore, I am indebted to Prof. Dr. Erez Berg for his willingness to co-examine my PhD and for his insightful comments. I would also like to mention the support I received from Dr. Frank Pollmann in the earlier stages of my PhD, when trying to set up the numerical machinery that played an important role later on.

The amazing atmosphere at the Institute for Theoretical Physics at ETH, in particular amongst the condensed matter groups, has provided a great incentive for working dilligently. I have had the pleasure to share an office with Murad Tovmasyan, Roman Süsstrunk and occasionally with Sina Zeytinoglu. I would like to thank them in particular for the small breaks and interesting discussions we had on various topics in physics, but also those about daily life. Such interesting discussions were not limited to the office, of course, and I would like to thank all the former and current members of the condensed matter groups with whom I have had the pleasure to discuss. They are Adel Benlagra, Adrien Bouhon, Martin Buchacek, Jonathan Buhmann, Tomas Bzdusek, Matteo Biondi, Chiara Decaroli, Daniel Mueller, Sarah Etter, Wei Chen, Michael Ferguson, Patrick Haughian, Dima Ivanov, Maciek Koch-Janusz, Andrei Lebedev, Markus Legner, Yehua Liu (多谢你!), Jan Lowinksi, David Moeckli, Paolo Molignini, David Oehri, Luca Papariello, Diana Prychynenko, Chitra Ramasubramanian, Aline Ramires, Laurent Schaller, Sebastian Schmidt, Roland Willa and Tobias Wolf. I also thank PD Dima Geshkenbein for introducing me to climbing, Prof. Dr. Oded Zilberberg for introducing me to capoeira, and Prof. Dr. Manfred Sigrist and Prof. Dr. Gianni Blatter for their support, knowledge, and for creating the amazing atmosphere amongst the theoretical condensed matter groups.

My utmost gratitude goes to the best, craziest and most important person in my whole universe, Rodinde. Thank you so much for believing in me, and for all the support you have given me over the past decade! Thank you for your understanding in stressful times. Thank you for being my *home*. Ik hou van je liefie, en ik kijk ook uit naar al onze volgende avonturen!

Lieve Pa & Ma, zonder jullie steun en vertrouwen zou ik dit nooit bereikt hebben! Be-

dankt voor alle hulp, alle boterhammen tijdens de middelbare school en alle filosofische discussies tot midden in de nacht! Bedankt voor al het begrip rond de verhuizingen naar Leiden en naar het buitenland, en voor het langskomen in Zwitserland!

Lieve Lodewijk & Michelle, bedankt voor alle koffie en de gezellige visite als we weer eens in Nederland waren! En lieve Aiden, ook jij bedankt voor alle gekke bekken en vrolijkheid tijdens die visite!

Lieve Harrie & Florentine, ik wil ook jullie heel erg bedanken voor alles wat jullie voor Rodinde en mij hebben gedaan! Alle etentjes, de vakanties, de musicals, ik zal ze nooit vergeten. Bedankt voor het met interesse volgen van al onze gebeurtenissen!

Lieve Faar, dankjewel voor alle gezelligheid als we op bezoek waren!

Lieve Quinnie, ook jij bedankt voor alle gezelligheid en je gastvrijheid tijdens onze Nederland-bezoeken. Ik hoop dat we nog vaak met z'n allen op vakantie kunnen gaan, zoals we met de treinreis en in Schotland gedaan hebben!