

University of Trento
Department of Physics,
University of Cologne
Faculty of Mathematics and Natural Sciences
Institute of Theoretical Physics

Data driven approach to detection of quantum phase transitions

Daniele Contessi

July 2023

Supervised by Dr. Alessio Recati, Prof. Elisa Ricci and Prof.
Matteo Rizzi

Submitted in partial fulfillment of the requirements for the degree of
Dottore di Ricerca in Fisica at the University of Trento and
Doktor der Naturwissenschaften at the University of Cologne
as stated in the Co-tutelle agreement of April 9th, 2021.

Abstract

Phase transitions are fundamental phenomena in (quantum) many-body systems. They are associated with changes in the macroscopic physical properties of the system in response to the alteration in the conditions controlled by one or more parameters, like temperature or coupling constants. Quantum phase transitions are particularly intriguing as they reveal new insights into the fundamental nature of matter and the laws of physics. The study of phase transitions in such systems is crucial in aiding our understanding of how materials behave in extreme conditions, which are difficult to replicate in laboratory, and also the behavior of exotic states of matter with unique and potentially useful properties like superconductors and superfluids. Moreover, this understanding has other practical applications and can lead to the development of new materials with specific properties or more efficient technologies, such as quantum computers. Hence, detecting the transition point from one phase of matter to another and constructing the corresponding phase diagram is of great importance for examining many-body systems and predicting their response to external perturbations.

Traditionally, phase transitions have been identified either through analytical methods like mean field theory or numerical simulations. The pinpointing of the critical value normally involves the measure of specific quantities such as local observables, correlation functions, energy gaps, etc. reflecting the changes in the physics through the transition. However, the latter approach requires prior knowledge of the system to calculate the order parameter of the transition, which is uniquely associated to its universality class. Recently, another method has gained more and more attention in the physics community. By using raw and very general representative data of the system, one can resort to machine learning techniques to distinguish among patterns within the data belonging to different phases. The relevance of these techniques is rooted in the ability of a properly trained machine to efficiently process complex data for the sake of pursuing classification tasks, pattern recognition, generating brand new data and even developing decision processes.

The aim of this thesis is to explore phase transitions from this new and promising data-centric perspective. On the one hand, our work is focused on the development of new machine learning architectures using state-of-the-art and interpretable models. On the other hand, we are interested in the study of the

various possible data which can be fed to the artificial intelligence model for the mapping of a quantum many-body system phase diagram. Our analysis is supported by numerical examples obtained via matrix-product-states (MPS) simulations for several one-dimensional zero-temperature systems on a lattice such as the XXZ model, the Extended Bose-Hubbard model (EBH) and the two-species Bose Hubbard model (BH2S).

In Part I, we provide a general introduction to the background concepts for the understanding of the physics and the numerical methods used for the simulations and the analysis with deep learning. In Part II, we first present the models of the quantum many-body systems that we study. Then, we discuss the machine learning protocol to identify phase transitions, namely anomaly detection technique, that involves the training of a model on a dataset of *normal* behavior and use it to recognize deviations from this behavior on test data. The latter can be applied for our purpose by training in a known phase so that, at test-time, all the other phases of the system are marked as anomalies. Our method is based on Generative Adversarial Networks (GANs) and improves the networks adopted by the previous works in the literature for the anomaly detection scheme taking advantage of the adversarial training procedure. Specifically, we train the GAN on a dataset composed of bipartite entanglement spectra (ES) obtained from Tensor Network simulations for the three aforementioned quantum systems. We focus our study on the detection of the elusive Berezinskii-Kosterlitz-Thouless (BKT) transition that have been object of intense theoretical and experimental studies since its first prediction for the classical two-dimensional XY model. The absence of an explicit symmetry breaking and its gappless-to-gapped nature which characterize such a transition make the latter very subtle to be detected, hence providing a challenging testing ground for the machine-driven method. We train the GAN architecture on the ES data in the gappless side of BKT transition and we show that the GAN is able to automatically distinguish between data from the same phase and beyond the BKT. The protocol that we develop is not supposed to become a substitute to the traditional methods for the phase transitions detection but allows to obtain a qualitative map of a phase diagram with almost no prior knowledge about the nature and the arrangement of the phases – in this sense we refer to it as agnostic – in an automatic fashion. Furthermore, it is very general and it can be applied in principle to all kind of representative data of the system coming both from experiments and numerics, as long as they have different patterns (even hidden to the eye) in different phases.

Since the kind of data is crucially linked with the success of the detection, together with the ES we investigate another candidate: the probability density function (PDF) of a globally U(1) conserved charge in an extensive sub-portion of the system. The full PDF is one of the possible reductions of the ES which is known to exhibit relations and degeneracies reflecting very peculiar aspects of the physics and the symmetries of the system. Its patterns are often used to tell different kinds of phases apart and embed information about non-local quantum correlations. However, the PDF is measurable, e.g. in quantum gas microscopes experiments, and it is quite general so that it can be considered not only in the cases of the study but also in other systems with different symmetries and dimensionalities. Both the ES and the PDF can be extracted from the simulation of the ground state by dividing the one-dimensional chain into two complementary subportions. For the EBH we calculate the PDF of the bosonic occupation number in a wide range of values of the couplings and we are able to reproduce the very rich phase diagram containing several phases (superfluid, Mott insulator, charge density wave, phase separation of supersolid and superfluid and the topological Haldane insulator) just with an educated gaussian fit of the PDF. Even without resorting to machine learning, this analysis is instrumental to show the importance of the experimentally accessible PDF for the task. Moreover, we highlight some of its properties according to the gapless and gapped nature of the ground state which require a further investigation and extension beyond zero-temperature regimes and one-dimensional systems.

The last chapter of the results contains the description of another architecture, namely the Concrete Autoencoder (CAE) which can be used for detecting phase transitions with the anomaly detection scheme while being able to automatically learn what the most relevant components of the input data are. We show that the CAE can recognize the important eigenvalues out of the entire ES for the EBH model in order to characterize the gapless phase. Therefore the latter architecture can be used to provide not only a more compact version of the input data (dimensionality reduction) – which can improve the training – but also some meaningful insights in the spirit of machine learning interpretability.

In conclusion, in this thesis we describe two advances in the solution to the problem of phase recognition in quantum many-body systems. On one side, we improve the literature standard anomaly detection protocol for an auto-

matic and agnostic identification of the phases by employing the GAN network. Moreover, we implement and test an explainable model which can make the interpretation of the results easier. On the other side we put the focus on the PDF as a new candidate quantity for the scope of discerning phases of matter. We show that it contains a lot of information about the many-body state being very general and experimentally accessible.

Contents

Abstract	i
I Background	1
1 Many-body states and quantum entanglement	3
1.1 Quantum Entanglement	4
1.2 Entropy (Von Neumann)	6
1.3 Area-law	8
1.4 Schmidt decomposition	10
1.5 Entanglement Spectrum	11
2 Numerical simulation of quantum many-body systems	13
2.1 Matrix Product States	18
2.2 Gauge freedom	20
2.3 Matrix product operators	22
2.4 MPO representation of the Hamiltonian	24
2.5 Variational iterative ground-state searching	26
2.6 Numerical Efficiency	29
2.7 Symmetries in tensor networks	30
2.8 Entanglement and correlations	32
2.9 Boundary conditions	32

3	Deep Learning	35
3.1	Structure and training of a NN: concepts and a concrete example	36
3.2	Autoencoders	44
4	Introduction to quantum phase transitions	49
II	Results	53
5	Models	57
5.1	XXZ model	57
5.2	Extended Bose Hubbard model	58
5.3	Two-component Bose-Hubbard model	60
5.4	Numerical details	63
6	Detection of the BKT transition with GAN	65
6.1	Anomaly detection scheme	65
6.2	GAN architecture	66
6.3	Features of the symmetry resolved ES	69
6.4	Further data analysis of the ES	73
6.5	Results with GAN	77
6.6	Loss threshold and precise transition detection	84
6.7	Conclusions	86
7	Gaussian fit of PDF	89
7.1	Number probability density function in the EBH model	91
7.2	PDF for gapped phases	95
7.3	Conclusions	99

8	Feature selection	101
8.1	Training for the ES dataset	104
8.2	Results for 256 sites BH model	107
8.3	Conclusions	111
9	Discussion and Conclusions	113
A	Luttinger parameters	115
	References	118
	Acknowledgements	130
B	Acknowledgements	131

Part I

Background

Chapter 1

Many-body states and quantum entanglement

A quantum state is the mathematical description of a quantum system which provides the probability distribution of the outcomes of possible measurements. When such a state is a vector of the Hilbert space $\psi \in \mathcal{H}$, the state is said to be pure. The expectation value of a generic operator \hat{A} is then:

$$\langle \hat{A} \rangle = \text{tr}(\rho \hat{A}) = \langle \psi | \hat{A} | \psi \rangle, \quad (1.1)$$

where $\rho = |\psi\rangle\langle\psi|$ is the density matrix operator that is positive semi-definite, Hermitian and has trace one. If the state is not pure but can be written as a linear combination of pure states, it is mixed. Usually mixed states are associated to fluctuations or errors in experiments because the system, when measured, will collapse in one of the pure states with a probability given by the coefficients of the linear combination. As a consequence of the spectral theorem, it always exists a set $\{|\phi_\lambda\rangle\}$ living in \mathcal{H} , so that the density matrix is decomposed:

$$\rho = \sum_{\alpha} p_{\lambda} |\phi_{\lambda}\rangle\langle\phi_{\lambda}|, \quad (1.2)$$

with the non-negative coefficients p_{λ} summing up to one. The expectation value of \hat{A} can then be written as:

$$\langle \hat{A} \rangle = \text{tr}(\rho \hat{A}) = \sum_{\lambda} p_{\lambda} \langle \phi_{\lambda} | \hat{A} | \phi_{\lambda} \rangle, \quad (1.3)$$

so the dependence on the probability distribution p_λ is explicit. While the trace of the density matrix is one, the trace of the squared density matrix takes the name of *purity*:

$$\text{tr}(\rho^2) = \sum_{\lambda} p_{\lambda}^2 \leq 1. \quad (1.4)$$

This quantity is considered as a measure of how much the state is pure. If it is one, then the distribution of p_λ is peaked on a single pure state, while if the distribution is completely flat meaning the probabilities are $p_\lambda = 1/\dim(\mathcal{H})$ the state is maximally mixed.

1.1 Quantum Entanglement

The concept of quantum entanglement is one of the most interesting and mysterious aspects of quantum mechanics that has gained increasing attention after the Einstein-Podolsky-Rosen (EPR) paradox formulation [1]. It is related to the correlations among the degrees of freedom of a quantum system which cannot be explained classically. Within the quantum information framework, entanglement is considered a resource for quantum tasks such as quantum computation and quantum communication. More specifically for our purpose, it is a universal compression that was proven to be very useful for the characterization of a quantum many-body system whose description is cursed by an exponentially growing number of degrees of freedom.

We can introduce entanglement through the quantification of the amount of information that can be retrieved from one system A about another system B which has been in contact with A. Following the discussion of Bell in its seminal work [2], A and B could be spin one-half particles created somehow in the singlet spin state – hence a pure state – and moving freely in opposite directions. The EPR paradox arises when we assume that *locality* is obeyed: when the A spin component is measured, for instance, in the z -direction and results $+1$, then according to quantum mechanics the measure of B spin z -component when it is far apart will be -1 . Assuming locality, the only way of explaining the correlated outcomes is that they must be predetermined. However, such predetermination is not contemplated in the quantum mechanical description of the singlet. One option to include predetermination of the result is to include an additional hidden variable in the quantum mechanical description

of the initial state as is formalized in [2] in form of one of the famous Bell's inequalities. Those were proven to be violated by quantum systems afterwards [3, 4, 5, 6]. Eventually, this “spooky action at a distance” phenomenon that violates locality took the name of entanglement.

In the following we consider the division of a pure system in two complementary subportions (bipartition) in a more systematic way. We denote with A and B two complementary subsystems with their respective Hilbert spaces \mathcal{H}_A and \mathcal{H}_B . Their union is AB with \mathcal{H}_{AB} its Hilbert space. When the state of the entire system takes the form of:

$$|\psi\rangle = |\phi^{(A)}\rangle \otimes |\chi^{(B)}\rangle \quad (1.5)$$

where $|\phi^{(A)}\rangle$ ($|\chi^{(B)}\rangle$) is a state in $\mathcal{H}_{A(B)}$ it is known as *separable* or *product state*. The negative definition of entanglement is: A is said to be entangled with B (or AB is entangled) when AB is **not** separable.

We can extend to the more general mixed case by retrieving the density matrices. The state of the entire system is separable when its density matrix can be written as:

$$\rho = \sum_{\lambda} p_{\lambda} \rho_{\lambda}^{(A)} \otimes \rho_{\lambda}^{(B)}, \quad (1.6)$$

where in this case $\rho_{\lambda}^{(\alpha)} = |\phi_{\lambda}^{(\alpha)}\rangle\langle\phi_{\lambda}^{(\alpha)}|$ for the α -subsystem and $\{|\phi_{\lambda}\rangle^{(\alpha)}\}$ a basis for \mathcal{H}_{α} .

Before treating the measures of entanglement, it is useful to relate it to the concept of mixedness. The reduced density matrix of one subsystem with respect to the other is computed with a partial trace on the total density matrix:

$$\rho_A^{\text{red}} = \text{Tr}_B[\rho] = \sum_i \langle\chi_i^{(B)}|\rho|\chi_i^{(B)}\rangle, \quad (1.7)$$

being $|\chi_i^{(B)}\rangle$ a basis for \mathcal{H}_B . Now if we consider a pure entire AB state, then the reduced density matrix of A is:

$$\rho_A^{\text{red}} = \sum_{\lambda} p_{\lambda} \rho_{\lambda}^{(A)}. \quad (1.8)$$

It is immediate to see that the condition for which the state is separable, so not entangled, implies that the distribution p_{λ} is peaked on a single vector of

the A Hilbert space:

$$\rho_A^{\text{red}} = |\phi^{(A)}\rangle \langle \phi^{(A)}|, \quad (1.9)$$

hence a pure state is said to be entangled when the reduced density matrix of one subsystem is not (classically) mixed. It is also straightforward to define a proper measure for entanglement of pure states starting, for instance, from the purity. The situation for mixed states is quite more involved and still cutting-edge of current research.

1.2 Entropy (Von Neumann)

The von Neumann entropy [7] is a straightforward adaption of the Shannon entropy [8] of a discrete probability distribution from information theory to density matrices in quantum statistics. Entropy is generally regarded as a measure for the average amount of information that we gain about a random variable by triggering and observing one outcome. The definition of the Von Neumann entropy is:

$$S(\rho) = -\text{Tr}[\rho \ln \rho] \quad (1.10)$$

where the trace operation is performed on a complete basis of the Hilbert space where the system lives. Through the decomposition of the density on a basis, the entropy has a more simple expression:

$$S(\rho) = -\text{Tr} \left[\sum_{\lambda} p_{\lambda} |\phi_{\lambda}\rangle \langle \phi_{\lambda}| \ln \left(\sum_k p_k |\phi_k\rangle \langle \phi_k| \right) \right] = -\sum_{\lambda} p_{\lambda} \ln[p_{\lambda}]. \quad (1.11)$$

If the system is in a pure state, the density matrix is idempotent and $S(\rho) = 2S(\rho)$ would hold, leading to the unique option of zero entropy. It can be further demonstrated that entropy is zero if and only if the state is pure while is maximum for a maximally mixed state, once p_{λ} is flat. One can use the Von Neumann entropy as a measure of the amount of mixedness of the reduced density matrix of the a bipartition of a pure state [9] and hence of the amount of entanglement $S(A) = -\text{Tr} \rho^{\text{red}} \ln \rho^{\text{red}}$. Indeed, other properties contribute to the eligibility of the Von Neumann entropy as a *measure of entanglement* [10, 11]. We expect a measure of entanglement to quantify the amount of information that can be retrieved from one subsystem about the other, but meanwhile it should not change if we perform a local operation on either one of the subsystems. This is the case for the Von Neumann entropy

which is invariant under the so-called Local Operations and Classical Communications (LOCC) which also account for classical broadcasting of information. For any valid reduced density matrix ρ^{red} – namely positive semi-definite, hermitian, trace-one operator – the entropy is real and non-negative. Moreover, it possesses the property of additivity: if A and A^* are independent systems, then $S(\rho_A \otimes \rho_{A^*}) = S(\rho_A) + S(\rho_{A^*})$. The collection of these properties can be considered as a recipe for a good measure of entanglement for bipartite pure states, i.e. a function giving a unique ordering of entangled states.

The situation for mixed states is rather more complicated. For instance, the bipartite Von Neumann entropy is not a well-defined measure in the latter case. One can easily convince oneself with a simple contradiction for the mixed, separable ensemble of two qubits A and B $\rho = \frac{1}{2}(|00\rangle\langle 00| + |11\rangle\langle 11|)$ whose Von Neumann entropy is $S(\rho_A^{\text{red}}) = \ln(2)$ due to the classical correlations. At this point, it is possible to reduce the entropy using LOCCs. Imagine that the two qubits are separated in space and in one laboratory Alice performs the following operations on the A qubit: either she project the qubit in the $|1\rangle$ state or she project it on $|0\rangle$ and then she applies σ_x . Once she has done one of the two operations, she communicates with Bob so he performs the same one on qubit B . The resulting composite state is described by $\rho' = |11\rangle\langle 11|$ which has zero entropy. This is clearly a contradiction because LOCCs are not supposed to change entanglement.

The object of our interest are pure ground-states of quantum many-body systems, so we will not focus on the thorny issue of mixed states. We mention some possibilities for the measure of entanglement for those states following the much more complete treatment in [11]:

- Entanglement of formation considers all the possible decompositions that realize a given mixed state in terms of pure states. The entanglement of formation is defined as the minimum weighted average (the weights are p_λ) of the bipartite von Neumann entropy of the pure states over all the possible realizations.
- Distance-based measures (or relative entropy) employs the distance between the mixed state and the nearest separable state. The measure is defined choosing a proper distance function between the density matrices (e.g. Kullback-Leibler, Uhlmann fidelity, etc.).

- Operational-based measures like *distillable entanglement* that considers the amount of Bell states which can be distilled from n -copies of the mixed density matrix by means of LOCCs; *entanglement cost* on the contrary counts the amount of maximally entangled states that are needed to realize n -copies of the mixed state.

1.3 Area-law

In our work, we are interested in the properties of the ground-state of one-dimensional quantum many-body system on a lattice. Whereas the general description of the quantum many-body state is very hard because of the exponential growth of the Hilbert space in the number of constituents, we are focusing on a very peculiar corner of it. In such a region the number of parameters for a proper description is not affected by the curse of dimensionality but is limited when the system is not critical or is polynomial for the critical one [12]. The concept of bipartition that was introduced so far for the practical quantification of the amount of entanglement, becomes instrumental in the formalization of the above statement. The limited growth of degrees of freedom in one-dimensional systems is indeed reflected in a linear increase of the leading term of the entanglement entropy in the size of the boundary between the bipartitions.




Figure 1.1: Sketch of a system bipartition in the two complementary subportions A and B with ∂A the boundary of the A part. We mainly deal with those states of the system which lie in a small corner of the entire Hilbert space and possess the area-law property of the entanglement entropy.

This phenomenon is actually more general, it is valid also for higher dimensional frameworks and is usually denoted as *area-law*. In formula:

$$S(\rho_\alpha^{\text{red}}) \propto L_\alpha^{d-1} \quad (1.12)$$

where L_α is the linear size of the d -dimensional subportion α . In the context of quantum field theory, this behaviour has some tradition started by the intriguing phenomenon of Bekenstein-Hawking black hole entropy [13, 14]. It was proved that this peculiar scaling happens for the free Klein-Gordon field [15] and for conformal systems [16] which could be related

in some cases to the physics of black holes. On the other side, the limited number of parameters for an efficient description of area-law systems, gained increased interest in the context of numerical simulation as will be clear from the next section.

A demonstration of the validity of area-law for one-dimensional gapped systems with a local hamiltonian can be found in [17] meaning that the entanglement entropy is bounded by a constant which is independent on the system's size. For gapless (critical) systems this is no longer the case and the scaling of the entropy versus the size is logarithmic. Moreover, due to the link with conformal field theory at criticality, the coefficient of the logarithm has to do with the central charge of the underlying conformal field theory [12, 16, 18]. Another noticeable subject is the form of the corrections to area-law – aka the sub-leading terms – which acquire a universal and peculiar behaviour in presence of topological order [19, 20].

At last, it is worth mentioning another perspective about the restricted number of parameters necessary for the treatment of the systems described before that act as a springboard for the discussion about the numerical method in the next section. The fact that a system can be simulated by considering only a relevant smaller portion of the Hilbert space is the key on which renormalization methods rely. We can easily introduce their operating principle in an intuitive way: imagine that a system living in a Hilbert space of a given dimension is described by a certain microscopic theory. In order to build an easier, effective description, all the non-relevant degrees of freedom must be integrated out while the original couplings in the Hamiltonian are modified accordingly (renormalized). If a limited number of parameters is sufficient to describe the system, the new effective theory will be a good approximation of the original one. The theory about renormalization procedure provides a systematic approach to this task and allows for a physically motivated cancellation of the irrelevant couplings. In general it can keep track of the coupling modification as the renormalization flow proceeds. Usually this methods are successful when there is a natural *separation of scales* involving distances/energies. The relevant phenomena for this thesis involve for example a diverging correlation length at criticality meaning that the low-energy scale (low-wavelength behaviour) dominates the physical properties of the system.

Since the system is one-dimensional, the interactions are local and we pursue the ground state properties, the scaling of entanglement as genuine quantum

correlations is anyway bounded at most to a logarithmic growth. Being true that it is still growing with the system size, the increase is not as drastic as for volume-law ($S(\rho_\alpha^{\text{red}}) \propto L_\alpha^d$). Moreover, when the system is not critical, area-law behaviour dominates and hence one can expect that a class of ansätze which is efficient in approximating low-entangled state (as generated by the Density Matrix Renormalization Group) provides a good description of the relevant states.

1.4 Schmidt decomposition

A lot of entanglement properties for a pure bipartite state can be explored starting from the *Schmidt decomposition*. As we commented above, a pure state is not always separable, therefore one should provide two orthonormal basis for the two complementary subspaces $|i\rangle_A$ and $|j\rangle_B$ to properly describe the total state as:

$$|\psi\rangle = \sum_{i=1}^{\dim(A)} \sum_{j=1}^{\dim(B)} \Lambda_{ij} |i\rangle_A |j\rangle_B, \quad (1.13)$$

where the matrix of the coefficients Λ_{ij} is generically a rectangular matrix since we do not assume that the two subspaces have the same dimension. One can always perform a singular value decomposition (SVD) of the matrix Λ which allows to write the rectangular matrix as the product $\Lambda = UDV^\dagger$ (see Fig. 1.2 for an illustration), where U and V are $\dim(A) \times \min(\dim(A), \dim(B))$ dimensional and $\min(\dim(A), \dim(B)) \times \dim(B)$ dimensional matrices with orthonormal columns and row respectively (one of them being squared matrix is unitary) and D is a diagonal $\min(\dim(A), \dim(B)) \times \min(\dim(A), \dim(B))$ matrix with the non-negative *singular values* λ_α . By applying the SVD, the state of Eq. (1.13) can be re-written as:

$$\begin{aligned} |\psi\rangle &= \sum_{i=1}^{\dim(A)} \sum_{j=1}^{\dim(B)} \sum_{\alpha=1}^{\min(\dim(A), \dim(B))} U_{i\alpha} \lambda_\alpha V_{\alpha j}^\dagger |i\rangle |j\rangle \\ &= \sum_{\alpha=1}^{\min(\dim(A), \dim(B))} \lambda_\alpha \left(\sum_{i=1}^{\dim(A)} U_{i\alpha} |i\rangle \right) \left(\sum_{j=1}^{\dim(B)} V_{\alpha j}^\dagger |j\rangle \right), \end{aligned} \quad (1.14)$$




Figure 1.2: Pictorial representation of the singular value decomposition of a rectangular matrix Λ . The matrix Λ is decomposed into the product of three matrices: U , D and V^\dagger . The columns of U and V^\dagger are orthonormal and the diagonal elements of D are the singular values λ_α .

and the Schmidt decomposition finally reads:

$$|\psi\rangle = \sum_{\alpha=1}^r \lambda_\alpha |\varphi_\alpha\rangle |\theta_\alpha\rangle, \quad (1.15)$$

where $|\varphi_\alpha\rangle$ and $|\theta_\alpha\rangle$ are orthonormal bases of A and B because of the orthonormality properties of U and V and $r = \min(\dim(A), \dim(B))$ is the Schmidt rank. The Schmidt decomposition provides a nice and numerically efficient way to compute the spectrum of the reduced density matrix because the partial trace can be easily applied to the state as written in (1.15):

$$\rho_A^{\text{red}} = \sum_{\alpha_1}^r \lambda_{\alpha_1}^2 |\varphi_{\alpha_1}\rangle \langle \varphi_{\alpha_1}|, \quad \rho_B^{\text{red}} = \sum_{\alpha_1}^r \lambda_{\alpha_1}^2 |\theta_{\alpha_1}\rangle \langle \theta_{\alpha_1}|. \quad (1.16)$$

Notice that the reduced density matrices of the two subspaces share the same spectrum with a number of eigenvalues bounded by the smaller dimension among the spaces. The singular values λ_α are also called *Schmidt values* and their square correspond to the eigenvalues of the reduced density matrices.

1.5 Entanglement Spectrum

With the tools introduced above, we can now define the bipartite *entanglement spectrum* (ES) of a pure state. The latter corresponds to the set of the logarithm of reduced density matrix eigenvalues of the two subspaces:

$$\xi_\alpha = -\log \lambda_\alpha^2. \quad (1.17)$$

so that the reduced density matrix of A , being always positive semi-definite, can be cast in an exponential form in the spirit of a partition function:

$$\rho_A^{\text{red}} = \frac{1}{Z} e^{-H_E} \quad \text{with} \quad H_E = \sum_{\alpha=1}^r \xi_{\alpha} |\varphi_{\alpha}\rangle \langle \varphi_{\alpha}| \quad (1.18)$$

where H_E is the *Entanglement Hamiltonian*. The constant Z , written in analogy to thermodynamics, ensures the correct normalization $\text{Tr}(\rho_A^{\text{red}}) = 1$ but can also be included in the exponent showing that the meaningful information in the ES is not contained in the absolute magnitudes of its eigenvalues but rather in their relative magnitudes. Despite its suggestive form which resembles the one of a system coupled with a thermal bath, the above expression is not a true Boltzmann formula because H_E does not correspond to the real Hamiltonian restricted to the A subsystem.

The ES and its Hamiltonian have been studied for a plethora of systems. The first achievements were obtained with the Bisognano-Wichmann theorem in the context of generic field theory [21] and for conformal field theories afterwards [22]. The seminal work by Li and Haldane [23] demonstrated that the low-lying part of the ES contains universal features associated with topological properties. This has aroused great interest and triggered the study of ES as a very powerful tool for the detection and characterisation of topological phases. Among numerous fascinating properties of the ES, the notion of entanglement Hamiltonian in its very natural definition has been also intensively debated and is still a research topic. A great summary of the obtained results about the ES can be found in [24].

Chapter 2

Numerical simulation of quantum many-body systems

As previously anticipated, the concepts grounded on the renormalization procedure had a huge impact also on the numerical simulation of quantum many-body systems. The idea is to build a good approximation of the ground state of the system by considering only a relevant portion of the Hilbert space. Several so-called *decimation procedures* for eliminating the irrelevant degrees of freedom for the proper description of the ground state of such systems began to be developed in the 1990s, following and adapting the original seminal work of Wilson [25, 26]. The translation of the original renormalization group procedures for the study of the thermodynamic-limit ground-state properties is based on the replacement of the energy levels with lattice sites. Unfortunately, the first schemes turned out to be very poor as shown in the work by White and Noack of 1992 [27] where a numerical real-space renormalization group was proven to be ineffective for a tight-binding model emulating the particle in a box. The procedure was anyway conceptually crucial for the further developments which lead to the introduction of the Density Matrix Renormalization Group (DMRG) [28, 29]. The first receipt for the real space numerical renormalization procedure is summarized as follows:

1. A sub-portion A of the system called *block* and containing L sites is chosen small enough so the Hamiltonian acting on it H_A can be exactly diagonalized
2. A compound block AA of length $2L$ is formed and the Hamiltonian

H_{AA} is formed accounting for the two single-blocks Hamiltonian and the interblock interactions.

3. H_{AA} is diagonalized by keeping only m lowest-lying eigenvalues.
4. H_{AA} is projected on the truncated space spanned by the m eigenvectors new matrix $H_{AA}^{\text{tr.}}$
5. Loop from step 2 until the whole system is covered by compound of blocks.

The key point is that the decimation (or truncation) of the highly energetic eigenstates of each block should not contribute to the ground state of the overall system. The numerical advantage lies in the low dimension of the Hamiltonian to be diagonalised after each decimation step. Unfortunately, even in very simple scenarios, the first procedure leads to some problems due to the boundary conditions of the single block. [27].

White's brilliant breakthrough [28] consisted in the observation that, in order to have a general procedure, one should carefully consider which states must be retained for a block A when it is immersed in an environment, i.e. the case of the true thermodynamic limit in which A is ultimately embedded. In this context, he showed that the set of states which is particularly suitable for representing the properties of a composite block AA for the decimation procedure is that formed by the leading eigenvectors of the reduced density matrix of AA . Assuming that a system composed of L sites with a Hilbert space of dimension $m^{(S)}$ spanned by a basis $\{|n_L^{(S)}\rangle\}$ is described by the Hamiltonian H_L . An *enlarged block* Hamiltonian H_{L+1} is created by adding a site to the system. The basis that spans the new space can be written as $\{|n_L^{(S)}\rangle \otimes |\sigma^{(S)}\rangle\}$ where $\{|\sigma^{(S)}\rangle\}$ is a local basis for the new site. Now the environment is mimicked by embedding the system in an environment of the same size by an analog construction, hence a *superblock* of length $2L + 2$ is formed. The approximation to the ground state at thermodynamic limit is chosen to be the ground state of the superblock, obtained by numerical diagonalization:

$$|\psi\rangle = \sum_{n_L^{(S)}=1}^{m^{(S)}} \sum_{\sigma^{(S)}=1}^d \sum_{\sigma^{(E)}=1}^d \sum_{n_L^{(E)}=1}^{m^{(E)}} \phi^{n_L^{(S)}, \sigma^{(S)}, \sigma^{(E)}, n_L^{(E)}} |n_L^{(S)}\rangle \otimes |\sigma^{(S)}\rangle \otimes |\sigma^{(E)}\rangle \otimes |n_L^{(E)}\rangle. \quad (2.1)$$

where the dimension of the system S (environment E) is $N^{(S)} = m^{(S)} \times d$ ($N^{(E)} = m^{(E)} \times d$). The decimation prescription which leads to the DMRG can be interpreted as one of the following three lines of argument:

- If the superblock state of (2.1) describes the physical state of the system through the reduced density matrix ρ_S^{red} (obtained by tracing out the environment) and we assume that its eigenvalues are ordered by magnitude $p_1 \geq p_2 \geq \dots \geq p_{N^{(S)}}$, then it can be demonstrated that retaining only the first $\chi^{(S)}$ dominant eigenvalues of ρ_S^{red} is the optimal choice for reducing the error in the expected value of a generic observable bounded on the system. The error in the estimate is proportional to the *truncated weight* $\sum_{\alpha=\chi^{(S)}+1}^{N^{(S)}} p_\alpha$.
- The quadratic norm of the distance $\| |\psi^{(S)}\rangle - |\tilde{\psi}^{(S)}\rangle \|$ of the state approximation $|\tilde{\psi}\rangle = \sum_{i=1}^{\chi^{(S)}} \sum_{j=1}^{N^{(E)}} \Lambda_{ij} |i\rangle |j\rangle$ with truncated space to the true generic state $|\psi\rangle = \sum_{i=1}^{N^{(S)}} \sum_{j=1}^{N^{(E)}} \Lambda_{ij} |i\rangle |j\rangle$ can be minimized by keeping the $\chi^{(S)}$ eigenvectors $|i\rangle$ of the system with the largest eigenvalues of the reduced density matrix ρ_S^{red} . The error is the truncated weight of the previous point.
- The entanglement-wise optimization can be applied after a Schmidt decomposition:

$$|\psi\rangle = \sum_{\alpha=1}^{N_{\text{Schmidt}}} \lambda_\alpha |w_\alpha^{(S)}\rangle |w_\alpha^{(E)}\rangle, \quad (2.2)$$

and the Schmidt number $N_{\text{Schmidt}} \leq \min(N^{(S)}, N^{(E)})$ once the assumption $N^{(S)} \geq N^{(E)}$ is relaxed. The truncation of the Schmidt values to $\chi^{(S)} \leq N_{\text{Schmidt}}$ states by keeping the largest in magnitude means that only a fraction of the entanglement between the system and the environment is allowed. The discarded, low-probable, correlations are those corresponding to the truncated values.

The formulation of the *infinite-system DMRG algorithm* [28] can now be described following the aforementioned procedure with the following steps:

1. Consider a lattice of L sites forming the system block S living in a Hilbert space spanned by $m^{(S)}$ states $\{|n_L^{(S)}\rangle\}$ and an Hamiltonian H_L^S bounded to the block. Another block for the environment E is formed accordingly.

2. The new system block S' is formed by adding a new site. The new basis is then $\{|n_L^{(S)}\rangle \otimes |\sigma^{(S)}\rangle\}$ containing $N^{(S)} = m^{(S)} \times d$ states. The same construction is performed for the environment $E \rightarrow E'$. The entire superblock of length $2L + 2$ contains S' and E' and the Hamiltonian acting on it is denoted by H_{2L+2} .
3. The Hamiltonian H_{2L+2} is diagonalized, its eigenvalues and eigenvectors are sorted in decreasing order.
4. The reduced density matrix of the system $\rho_{S'}^{\text{red}} = \text{Tr}_{E'} \rho$ is determined and only $\chi^{(S)}$ leading eigenstates are kept in order to form a new reduced basis for S' . The reduced diagonalizing matrix composed of the kept $\chi^{(S)}$ on the columns is denoted by T and is $N^{(S)} \times \chi^{(S)}$ dimensional.
5. The new effective (truncated) Hamiltonian for the system S' is formed by decomposition on the new reduced basis of the block $H_{L+1}^{\text{tr}} = T^\dagger H_{L+1}^S T$ and becomes the new Hamiltonian of the system as $S' \rightarrow S$ and $E' \rightarrow E$ for the next iteration starting from step 2. The iterations stop once the final wanted length for the system is reached.

During the DMRG steps, observables can be computed both at the end of the optimization and between every growing step. A detailed description of the observables' computation is provided later below, once the DMRG algorithm is re-formulated in the context of matrix product states.

The infinite DMRG method as in its original prescription leads to some problems mainly due to the fact that in the early steps the blocks are small and the chosen relevant states are likely to be very different from a block of the same size embedded in the bulk of the big final system. The *finite DMRG* manages to substantially cure these problems. The idea is that once the infinite DMRG steps lead to an L sized system, the growing procedure is stopped and an optimization of the obtained superblock is applied. The optimization recalls the same steps of the infinite algorithm but the growth of one block (e.g. the system) is accompanied by the shrinkage of the complementary block (e.g. the environment). The procedure is quite efficient if the infinite algorithm stages are stored so that for every optimization step, the description of the blocks is already available. This further processing of the final state is called sweeping: every sweep proceeds until the shrinking of the system or the environment reaches some minimum size and the growing can start again

at the expense of the complementary block. Convergence is usually checked through stabilization of the results from sweep to sweep. It is worth mentioning that convergence must be checked also varying the number of kept states $\chi^{(S)}$ as free parameter of the algorithm since it can happen that the apparently converged finite-system results are instead effect of the trapping in a low-entangled metastable state.

From now on, we will report the mechanisms of the DMRG in a more detailed way but by following a more modern path. The latter derives from a unrelated initial development of the optimization of a particular class of ansatz devoted to the approximation of quantum systems, namely the Matrix Product States (MPS). The early stages of the adoption of this family of variational states lead to several astonishing advances, one of them being the exact expression of the one-dimensional AKLT state [30]. The MPS connection to DMRG was done only a few years after their discovery [31, 32]. Eventually, it was recognized that the finite DMRG gives rise to quantum states which can be written in MPS form [33].

The DMRG algorithm and its expression in terms of MPS language, represented a fundamental milestone for a new state-of-the-art level of accuracy in the simulation of one-dimensional quantum many-body systems. Currently, for such systems, it still maintains this supremacy in the majority of the cases. These systems are usually hard to study and an exact solution is available only in few cases (e.g. with Bethe ansatz [34], exact diagonalization, quantum Monte Carlo [35]). In many applications of DMRG, the accuracy is limited by machine precision, quite independently of the nature of the microscopic Hamiltonian. In this work we are only concerned with the simulation of the ground state but the possibilities do not end here and extend to other purposes as real-time dynamics, excited states or thermal ensembles. Moreover, thanks to the very general map of a lattice system to a Tensor Network developed from the MPS, the idea could be extended to the simulation of two-dimensional systems by means of the projected entangled-pair states (PEPS) [36]. Despite the technical complexity of the algorithm and their computational demanding nature that make PEPS not able to reach the same standing as MPS yet, they are considered a promising tool for the simulation of two-dimensional systems and have already been used to study a wide range of physical phenomena, such as quantum phase transitions and topological phases of matter. Advances in computational resources and algorithmic developments are expected to over-

come some of their limitations in the future.

2.1 Matrix Product States

In the following, we will show that any one-dimensional pure quantum state can be written exactly in MPS form. Once this is established, the natural decimation scheme applicable to MPSs is presented so that they can be intuitively linked to the numerical procedure of DMRG. Being a very convenient language, supported also by an opportune pictorial representation, the MPS formalism is also used to present the detailed optimization procedure which leads to the determination of the ground state of one-dimensional quantum systems as well as the numerical calculation of observables (e.g. local observables, correlators, energy, etc.), of the entanglement properties and the realization of boundary conditions.




Figure 2.1: a pictorial representation of the $|\uparrow\downarrow\uparrow\rangle$ 3-qubits state. In particular the cube contains the coefficients of the tensor $\phi_{i_1 i_2 i_3}$ whose indices run over the single particle doublets.

We can introduce the MPS starting from a generic one-dimensional normalized pure quantum state on a lattice of L sites with d -dimensional local state spaces $\{i_j\}_{j=1,\dots,L}$ that can be written as:

$$|\psi\rangle = \sum_{i_1, i_2, \dots, i_N=1}^d \phi_{i_1 i_2 \dots i_L} |i_1\rangle \otimes |i_2\rangle \otimes \dots \otimes |i_L\rangle \quad (2.3)$$

where the coefficients $\phi_{i_1 i_2 \dots i_L}$ are exponentially many (d^L) in the lattice size and are collected in a big tensor structure (see Fig.2.1 for an example of a 3 spin-half qubits) where the dimensions of the tensor retrace the local physical ones. A

decomposition of such a tensor can be done in terms of smaller objects via additional new indices following a systematic procedure.

The idea is to separate the state in blocks by means of the SVD: we can start by reshaping the tensor into a matrix of dimension $d \times d^{L-1}$ that will be indicated with the notation $\phi_{i_1(i_2 \dots i_N)}$ after the *fusion* of all the indices but from the first one; the unique new index contains the Cartesian product of all the fused ones [37] (see Fig. 2.2 for the illustration of the fusion operation). The state




Figure 2.2: Coefficient of the many-body wavefunction in their tensor shape drawn in the tensor network diagrammatic representation (left). The open legs are the physical indices. The version of the tensor with fused links is depicted on the right. The new index j is the cartesian product of all the fused ones.

now reads:

$$|\psi\rangle = \sum_{i_1}^d \sum_{i_2 \dots i_N}^d \phi_{i_1(i_2 \dots i_N)} |i_1\rangle |i_2 \dots i_N\rangle. \quad (2.4)$$

We have already seen that SVD allows to separate the tensor in two contributions bounded to the complementary parts of the Hilbert space:

$$\phi_{i_1(i_2 \dots i_N)} = \sum_{\alpha_1}^{r_1} U_{\alpha_1}^{i_1} \lambda_{\alpha_1} (V^\dagger)_{\alpha_1}^{(i_2 \dots i_N)}. \quad (2.5)$$

Notice that the physical indices have been written as superscript just for the sake of clarity. From now on, we are not going to assign an additional index to tensors that contain the physical one i_l for indicating their position along the lattice since they are unambiguously identified by it; for instance $U_{\alpha_{l-1}\alpha_l}^{i_l} \equiv (U_l)_{\alpha_{l-1}\alpha_l}^{i_l}$. Reabsorbing the terms as done before, the Schmidt decomposition yields:

$$|\psi\rangle = \sum_{\alpha_1}^{r_1} \lambda_{\alpha_1} \left(\sum_{i_1}^d U_{\alpha_1}^{i_1} |i_1\rangle \right) \left(\sum_{i_2 \dots i_N}^d (V^\dagger)_{\alpha_1}^{(i_2 \dots i_N)} |i_2 \dots i_N\rangle \right) = \sum_{\alpha_1}^{r_1} \lambda_{\alpha_1} |\varphi_{\alpha_1}\rangle |\theta_{\alpha_1}\rangle, \quad (2.6)$$

The virtual index α_1 has been introduced running up to the Schmidt rank r_1 . The reduced density operators for the two subspaces – the first site and the rest of the system – can be written in the same spirit of Eq. (1.16).

The same protocol can be iteratively repeated for splitting the indices of all the other sites and resummings consecutively the diagonal matrices λ on the right after the SVDs decompositions. Ultimately, what is left is an exact decomposition in a (*left-canonical*) MPS:

$$|\psi\rangle = \sum_{i_1 \dots i_L} \sum_{\alpha_1 \dots \alpha_L} A_{\alpha_1}^{i_1} A_{\alpha_1 \alpha_2}^{i_2} \dots A_{\alpha_{L-2} \alpha_{L-1}}^{i_{L-1}} A_{\alpha_{L-1}}^{i_L} |i_1 \dots i_L\rangle. \quad (2.7)$$




Figure 2.3: Pictorial representation of the matrix product state. Every tensor (blue ball) is linked with the neighbours by closed legs meaning that the indices are contracted. The open legs (i_m) represent the physical indices running over the local Hilbert space basis.

It is called left-canonical because according to the recipe used for the decomposition, all the tensors have orthonormal columns $A_{\alpha_j \alpha_{j+1}}^{i_j} = U_{(\alpha_j) \alpha_{j+1}}^{i_j}$ for all $i \in \{1, L-2\}$ but for the last one $A_{\alpha_{L-1}}^{i_L} = \lambda_{\alpha_{L-1}} (V^\dagger)_{\alpha_{L-1}}^{i_L}$.


The decomposition does not change the order of the amount of coefficients if at each SVD the number of kept singular values is equal to the lowest of the dimensions of the decomposed matrix. Maximally, the dimensions of the A matrices is then $(1 \times d), (d \times d^2), \dots, (d^{L/2-1} \times d^{L/2}), (d^{L/2} \times d^{L/2-1}), \dots, (d \times d^2), (1 \times d)$ recaling the exponential growth of degrees of freedom. In order to make the MPS tractable from a numerical point of view, the dimension of the matrices must be truncated up to a number of allowed virtual states χ at the most. Therefore, from d^L coefficients, the number becomes of the order $L \times \chi^2 \times d$ that is polynomial in the dimension of the system. χ takes the name of *bond dimension* and is a free-parameter of the simulation which is strictly linked to the representational power of the MPS. We will see immediately how this parameter can be related to the truncated weight of DMRG and puts a limit on the maximal amount of entanglement present in the MPS state. In particular, the procedure for the decimation of the states in the DMRG language becomes a truncation of the kept singular values' maximum number.

The MPS structure can be very conveniently graphically represented as in Fig. 2.3 where every A matrix is a circular box, the physical indices are vertical legs and the α_j indices are horizontal legs connecting the matrices (all connected legs are summed over).

2.2 Gauge freedom

The MPS of Eq. (2.7) is written in general terms. The adopted sequential decomposition relying on the SVD that was carried out to transform a general state in the MPS form is not unique meaning that different MPSs can represent the same state. This redundancy is associated with a *gauge freedom* and

concretely corresponds to an invariance under a set of linear transformations of the tensors A of the MPS chain which have no impact on the physical degrees of freedom. In practice, it is easy to understand such a freedom considering the following illustration:



One can insert an identity $\mathbb{1} = X^{-1}X$ inside the closed leg between tensors A and B and reabsorb X^{-1} and X in A and B respectively, giving rise to two new tensors A' and B' . Since this operation is always possible, if the MPS provides a good description of a physical state, the physical information cannot be affected by this operation. The latter property can be instead exploited to drastically increase efficiency in the numerical simulation.

Nevertheless, let us recall the fact that the MPS of (2.7) is left-canonical: every tensor A but the last one on the right fulfills this property:

$$\sum_{\alpha} A_{\alpha\beta}^{i_m} A_{\alpha\beta'}^{\dagger i_m} = \mathbb{1}_{\beta\beta'} \quad (2.8)$$

that pictorially is shown in Fig. 2.4 where the left canonical tensors are indicated with a triangle so that they can be distinguished from the more general circular ones. By performing the same iterative decomposition as before but




Figure 2.4: In the picture the diagram of the equation (2.8). The shape of the tensor indicates the direction of the contraction used to employ the left canonical form obtained with the gauge transformation.

starting from the right, it is straightforward to obtain a *right-canonical* MPS for which the first tensor on the left now is written as $A_{\alpha_1}^{i_1} = U_{\alpha_1}^{i_1}$ and all the others $A_{\alpha_j\alpha_{j+1}}^{i_j} = (V^\dagger)_{\alpha_j\alpha_{j+1}}^{i_j}$. The equivalence of the two forms is just another example of the gauge freedom, it is indeed possible to transform a left-canonical MPS into a right-canonical one by absorbing the singular values matrix λ_α sequentially on the adjacent matrix U on the left and performing

again an SVD so that a sequence of V^\dagger is kept on the right and ultimately the first tensor becomes $A_{\alpha_1}^{i_1} = U_{\alpha_1}^{i_1}$. It is also possible to mix the decomposition of the state from left and right hence obtaining an MPS in general *canonical* form. The tensor of coefficients $\phi_{i_1 i_2 \dots i_L}$ is decomposed in the following way:

$$\phi_{i_1 i_2 \dots i_L} = \sum_{\alpha_1 \alpha_2 \dots \alpha_L} U_{\alpha_1}^{i_1} U_{\alpha_1 \alpha_2}^{i_2} \dots U_{\alpha_{l-1} \alpha_l}^{i_l} \lambda_{\alpha_l} V_{\alpha_l \alpha_{l+1}}^\dagger \dots V_{\alpha_{L-1} \alpha_L}^\dagger \quad (2.9)$$

The special site l is often called orthogonality center since all the tensors on the left and on the right possess orthogonality properties. Eventually, in Fig. 2.5 the reason why the gauge freedom is a resource for numerical calculations manifests clearly for the computation of the norm of a given MPS. The norm can be represented just as the contraction of all the physical legs of the MPS in Eq. (2.7) with the correspondent physical legs of the adjoint version of the same MPS. The operation is sketched in the Fig. 2.5 for a left-canonical MPS. Due to the orthonormality properties in this gauge, all the tensors but the far right one respect Eq. (2.8) and the contractions become trivial. The non-trivial contraction is the one between the last tensor and its adjoint which is the only one that must be computed. The procedure is usually employed also for the




Figure 2.5: Procedure for the contraction aimed at computing the norm of a state in the MPS representation. The definition of an orthogonality center allows to spare numerical effort. A smart employment of the gauge freedom leads to the most convenient form of the MPS in such a way that the norm consists in a unique double contraction (last step).

sake of computing the expectation values of local operators: the orthogonality center is placed on the site on which the local operator acts, thus simplifying the contractions on its left and right.

2.3 Matrix product operators


A natural generalization of the MPS structure can be written for operators, the idea is to consider the matrix elements of operators as written as chain of

tensors like for the state coefficients:

$$\langle \vec{i} | \hat{O} | \vec{i}' \rangle = W^{i_1, i'_1} W^{i_2, i'_2} \dots W^{i_{L-1}, i'_{L-1}} W^{i_L, i'_L} \quad (2.10)$$

where the matrices W are the analog of matrices A of (2.7), and $|\vec{i}\rangle = |i_1, \dots, i_L\rangle$ correspond to physical indices.

For carrying out the generalization explicitly, let us consider first the general shape of an operator acting on the many-body Hilbert space:



$$\hat{O} = \sum_{\{\vec{i}\}\{\vec{i}'\}} o_{\vec{i}, \vec{i}'} |\vec{i}\rangle \langle \vec{i}'| \quad (2.11)$$

Figure 2.6: Representation of the building block of an MPO.

The main difference from the state coefficients is that the tensor $o_{\vec{i}, \vec{i}'}$ has both ingoing and outgoing physical legs.

Any operator can actually be brought into the form of a chain of matrices by following the same idea of decomposition of the MPS, indeed $o_{\vec{i}, \vec{i}'} = o_{i_1 i_2 \dots i_L, i'_1 i'_2 \dots i'_L}$ can be reshaped as $o_{(i_1 i'_1)(i_2 i'_2) \dots (i_L i'_L)}$ and then decomposed in the same way as in the previous section. The resulting expression for the operator yields:

$$\hat{O} = \sum_{\{\vec{i}\}\{\vec{i}'\}} \sum_{\alpha_1, \dots, \alpha_L} W_{\alpha_1}^{i_1, i'_1} W_{\alpha_2, \alpha_3}^{i_2, i'_2} \dots W_{\alpha_{L-2}, \alpha_{L-1}}^{i_{L-1}, i'_{L-1}} W_{\alpha_L}^{i_L, i'_L} |\vec{i}\rangle \langle \vec{i}'|. \quad (2.12)$$

and takes the name of *Matrix Product Operator* (MPO). The building blocks of the MPO are tensors with 4 legs, two physical legs running on the input and output physical state's basis and the *MPO bond indices* (see Fig. 2.6). The application of a MPO to a MPS runs as:

$$\begin{aligned} \hat{O} |\psi\rangle &= \sum_{\{\vec{i}\}\{\vec{i}'\}} \sum_{\alpha_1, \dots, \alpha_L} \sum_{\beta_1, \dots, \beta_L} (W_{\alpha_1}^{i_1, i'_1} W_{\alpha_2, \alpha_3}^{i_2, i'_2} \dots W_{\alpha_{L-2}, \alpha_{L-1}}^{i_{L-1}, i'_{L-1}} W_{\alpha_L}^{i_L, i'_L}) \cdot \\ &\cdot (A_{\beta_1}^{i'_1} A_{\beta_1}^{i'_1} \dots A_{\beta_{L-2}, \beta_{L-1}}^{i'_{L-1}} A_{\beta_{L-1}}^{i'_L}) |i\rangle \\ &= \sum_{\{i\}} \sum_{\alpha_1, \dots, \alpha_L} \sum_{\beta_1, \dots, \beta_L} B_{\alpha_1 \beta_1}^{i_1} B_{\alpha_1 \alpha_2 \beta_1 \beta_2}^{i_2} \dots B_{\alpha_{L-1} \alpha_L \beta_{L-1} \beta_L}^{i_{L-1}} B_{\alpha_L \beta_L}^{i_L} |i\rangle \end{aligned} \quad (2.13)$$

where in the rightmost side the contractions $B_{\alpha_j \alpha_{j+1} \beta_j \beta_{j+1}}^{i_j} = \sum_{i'_j} (W_{\alpha_j, \alpha_{j+1}}^{i_j, i'_j} A_{\beta_j}^{i'_j})$ are performed. The result of the multiplication is still an MPS with a dimension corresponding to the combination of the original MPS and MPO dimen-

sions.

The computation of local observables expectation values becomes straightforward thanks to the MPO representation of operators. Moreover, because of the clear pictorial representation, the strategy for the contractions becomes very intuitive.

To make a concrete example, let's consider a local operator acting on the third site of the chain:

$$\begin{aligned}\hat{O}_3 &= \mathbb{1}_1 \otimes \mathbb{1}_2 \otimes \hat{O}_3 \otimes \mathbb{1}_4 \dots \mathbb{1}_L \\ &= \sum_{\{\vec{i}\}\{\vec{i}'\}} \mathbb{1}^{i_1 i'_1} \mathbb{1}^{i_2 i'_2} o_3^{i_3 i'_3} \mathbb{1}^{i_4 i'_4} \dots \mathbb{1}^{i_L i'_L} |\vec{i}\rangle \langle \vec{i}'| = \sum_{i_3, i'_3} o_3^{i_3 i'_3} |i_3\rangle \langle i'_3|\end{aligned}\quad (2.14)$$

is already in MPO form. In order to simplify the calculations one can exploit the gauge freedom in such a way that the minimum number of contractions is performed. The point is to properly define the orthogonality center of the canonical form that is the tensor for which on the right(left) everything is in the right(left) canonical form. Also in this case, the expectation value $\langle \psi | \hat{O}_3 | \psi \rangle$ on the state ψ written in MPS form is conveniently represented with the illustration in Fig. 2.7.




Figure 2.7: Simplification of the contractions for the expectation value of the local operator \hat{O}_3 . The orthogonality center is taken to be the third site.

2.4 MPO representation of the Hamiltonian

In this section we will cover the procedure for the translation of the Hamiltonian in MPO language. We do it for a concrete simple example as the transverse field Ising model Hamiltonian in order to grasp the principles. Generalizing to other models has a quite straightforward implementation though. The Hamiltonian involves only on-site and nearest-neighbour term, it reads:

$$H = \sum_j -J \hat{S}_j^x \hat{S}_{j+1}^x + h \hat{S}_j^z \quad (2.15)$$

where $S^{x,z}$ are the components of the spin operator labeled with the site's number of the lattice. The full Hamiltonian in its expanded version corresponds to the sum of locally acting operators involving generic neighbouring sites l and $l + 1$:

$$H_l = -J_l(\mathbb{1} \otimes \mathbb{1} \otimes \dots \otimes \mathbb{1} \otimes \hat{S}_l^x \otimes \hat{S}_{l+1}^x \otimes \mathbb{1} \otimes \dots \otimes \mathbb{1}) + h_l(\mathbb{1} \otimes \mathbb{1} \otimes \dots \otimes \mathbb{1} \otimes \hat{S}_l^z \otimes \mathbb{1} \otimes \dots \otimes \mathbb{1}). \quad (2.16)$$

The translation procedure to the MPO form retraces the so called *matrix product diagrams* or finite automata/state-machine [38, 39, 40]. We consider a local building block of an MPO, namely $W_{\alpha_l, \alpha_{l+1}}^{i_l, i_{l+1}}$, not as different matrices for each physical index, it is instead promoted to an operator-valued matrix $\hat{W}_{\alpha_l, \alpha_{l+1}}$ acting only locally at site l where the bond indices run over virtual states. Furthermore, we adopt the convention that when the operator-valued




Figure 2.8: Drawing of the local operator in the finite state machine language. The initial state I and the final state O are connected with an intermediate one (A) that is reached after the appearance of the spin operator. Another spin operator along the x direction is applied in order to reach the output state. On the other side the external field along z . The identity operators applied to the initial and final state are intended to make the picture the most general possible, both for boundary operators and bulk ones.

matrices are multiplied together over the bond indices, the contained operators are combined through the outer-product. We aim to obtain the expression for each matrix so that the full combination along the chain gives the total Hamiltonian $H = \hat{W}_{\alpha_1} \hat{W}_{\alpha_1, \alpha_2} \dots \hat{W}_{\alpha_{L-1}, \alpha_L} \hat{W}_{\alpha_L}$. For the transverse field Ising model, the matrix is the following:

$$\hat{W}_{\alpha_l, \alpha_{l+1}} = \begin{matrix} & \text{I} & \text{A} & \text{O} \\ \text{I} & \hat{\mathbb{1}} & \hat{S}_l^x & h_l \hat{S}_l^z \\ \text{A} & 0 & 0 & -J_l \hat{S}_l^x \\ \text{O} & 0 & 0 & \hat{\mathbb{1}} \end{matrix} \quad (2.17)$$

where we introduced 3 virtual states (I,A and O) over which the bond indices α_l, α_{l+1} run. Placing one term inside the matrix means selecting the operator to be applied to account for the transition from one input state in the l bond to the output state in the following $l + 1$. One can think of running over the string of (2.16) from left to right: the string of identities that does not change the initial state called I is taken into account by filling the transition I \rightarrow I with the identity. Another possibility is to encounter the non-trivial operator \hat{S}^x that changes the virtual state from I \rightarrow A. After that, the subsequent factor is fixed to $-J\hat{S}^x$ so from A one goes to another different state O and the former factor sits in A \rightarrow O. The state O stands for a completed interaction term for the l bond and one will find only identities afterwards. O \rightarrow O is again trivially the identity operator. At this point, the missing term is the transverse field, since at I no non-trivial operator has already been found along the string and from O one can find only identities, the $h\hat{S}^z$ operator occupies the transition from I \rightarrow O, without any needed intermediate state. All the other transitions are forbidden. The intuitive sketch of the transition is reported in Fig. 2.8.

It is easy to prove that the contraction of two subsequent matrices following the aforementioned procedure $\hat{W}_{\alpha_l \alpha_{l+1}} \hat{W}_{\alpha_{l+1} \alpha_{l+2}}$ allows for the collection of the term $-J\hat{S}_l^x \hat{S}_{l+1}^x + h\hat{S}_l^z + h\hat{S}_{l+1}^z$ on the top right corner. If the Hamiltonian has open boundary conditions, the first and the last matrices will be a row vector (of operators) and a column vector respectively. The full contraction on the bond indices yields the total Hamiltonian. This representation is formally proven to be optimal and leads to a bond size of the MPO tensors of $\chi_l = N_l + 2$ where N_l is the number of interaction operators crossing the bond $(l, l + 1)$ [41].

2.5 Variational iterative ground-state searching

The algorithm for computing the approximation to the ground state is based on a variational technique on the class of MPS states. Suppose to have an initial MPS ansatz that approximates the state that we denote with $|\psi(\{A_i\})\rangle$ – this could be an MPS filled with random numbers of the desired shape and size. The goal is that of optimizing its parameters (its tensors A_i) so that the

energy is minimized ([37, 42]):

$$\begin{aligned} \min_{\{A_i\}_{i=1,\dots,L}} \frac{\langle \psi(\{A_i^*\}) | \hat{H} | \psi(\{A_i\}) \rangle}{\langle \psi | \psi \rangle} &= \min_{\{A_i\}, \lambda} (\langle \psi | \hat{H} | \psi \rangle - \lambda \langle \psi | \psi \rangle) \\ &= \min_{\lambda} \left[\frac{\partial}{\partial \{A_i^*\}} (\langle \psi | \hat{H} | \psi \rangle - \lambda \langle \psi | \psi \rangle) = 0 \right]. \end{aligned} \quad (2.18)$$

On the right side of the equation, the minimization is redefined as extremization with the introduction of the Lagrange multiplier λ . Eventually, the optimal $|\psi(\{A_i\})\rangle$ will be the ground state approximation and λ its energy. The total minimization should be done simultaneously on all the tensors making this a hard non-linear optimization problem. As an approximation, the minimum over λ can be performed iteratively one site at a time. To clarify, we explicitly write the extremization involving the A_i tensor:

$$\begin{aligned} \min_{\lambda} \left[\frac{\partial}{\partial \{A_i^*\}} (\langle \psi | \hat{H} | \psi \rangle - \lambda \langle \psi | \psi \rangle) = 0 \right]_{\forall i} &\Leftrightarrow \\ \frac{\partial}{\partial \{A_i^*\}} A_i^\dagger H^{\text{eff}} A_i = \lambda \frac{\partial}{\partial \{A_i^*\}} A_i^\dagger N^{\text{eff}} A_i &\quad \forall i, \end{aligned} \quad (2.19)$$

where $H^{\text{eff}} = H_{A_1, \dots, A_{i-1}, A_{i+1}, \dots, A_L}^{\text{eff}}$ is the contracted MPO of the Hamiltonian over all the tensors of the MPS except for the i -th one. N^{eff} is nothing else but the overlap $\langle \psi | \psi \rangle$ fully contracted except for A_i (see Fig. 2.11 for the schematic representation). The single optimization step does not lead to the optimal state but will lower the energy and find a variationally better approximation of the ground-state. Usually the iterations are stopped when the energy does not change anymore or when a maximum number of iterations is reached.

The derivative of a chain of tensors with respect to one them corresponds to the whole chain without that specific one. The intuitive diagram is in Fig. 2.9 below. Therefore, Eq. (2.19) is that of Fig. 2.10. The latter is in the form of




Figure 2.9: Pictorial representation of the derivative of a chain of tensors with respect to one of them.




Figure 2.10: Pictorial representation of the step of the variational iterative ground-state search.

a generalized eigenvalue problem

$$\min_{\lambda} \quad H^{\text{eff}} A_i = \lambda N A_i \quad \forall i, \quad (2.20)$$

for every site of the chain: in order to find the MPS approximation of the ground state of a given Hamiltonian, one starts from an initial unbiased ansatz and iteratively optimizes the tensors by solving a generalized eigenvalue problem for every site.

The numerical efficiency, can be boosted by exploiting once again the canonical form. The choice of the orthogonality center on the operator A^i to be optimized allows for a much simpler expression for the right hand side of (2.20) since all the tensors for the overlap on the left and on the right of A^i contract to the identity. In such a gauge, the generalized eigenvalue problem becomes a standard eigenvalue problem that can be solved more efficiently for the lowest eigenvalues and eigenvectors (see Fig. 2.11 for the diagram of the simplification and Fig. 2.12 for the new form of equation (2.20)). For this scope, the usually adopted technique involves Lanczos [43] or Jacobi-Davidson [44] large sparse matrix solvers to decrease the complexity when the dimension of the matrix to be diagonalized (at worst $d\chi^2$ with χ the maximal bond dimension) is too large for an exact full diagonalization. The standard procedure for the iterative optimization is to consecutively sweep on the chain from left to right and viceversa so that the orthogonality center of the MPS can be moved from one site to the neighbouring one minimizing the operations. It probably does not sound so surprising at this point that the aforementioned ground state searching can be related to the finite-size version of DMRG presented in the introduction of the chapter. Actually the two methods are exactly identical if the optimization of the DMRG superblock with the sweeping procedure is done on a single site and not on two sites (one of the system S and the other of the environment E) at the same time. On the contrary, the MPS optimization




Figure 2.11: Schematic representation of the simplification of the right hand side of the generalized eigenvalue problem (2.20) by exploitation of the canonical form.




Figure 2.12: Final simplified version of the generalized eigenvalue problem (2.20) in the canonical form.

can be mapped exactly to the finite DMRG by optimizing two neighbouring tensors on the chain simultaneously. A very nice description of the relation between the two languages can be found in Ref. [33]

2.6 Numerical Efficiency

The MPS gauge as well as the implementation of the iterations for the ground state searching have no unique prescription. The freedom of choosing the order of the different contractions – which are in the end just pairwise matrix multiplications – allows for the optimization of the numerical scaling. The number of operations for a contraction, indeed, strongly affects the final applicability of the method since a more costly scaling can prevent the simulation to run in reasonable time for a given system.

It is very easy to compute the number of operations needed during a contrac-

tion: we count the number of dimensions of all open indices of both tensors and we multiply by the dimensions of the indices that will be contracted over. For instance the contraction of two tensors with indices α, β and γ, β, δ where $\alpha, \gamma, \delta = 1, \dots, \chi$ and $\beta = 1, \dots, d$ over β requires $d\chi^3$ operations. The order of operations for the norm of the MPS of Fig. 2.3 with a bond dimension χ and a physical dimension d without resorting to the canonical form is $\mathcal{O}(d\chi^3L)$, hence linear in the system size. This represents just one of the examples where the smart employment of the canonical form can drastically reduce the computational cost.

2.7 Symmetries in tensor networks

When dealing with numerical efficiency, one essential aspect is represented by the embedding of symmetries in the tensor network formalism. From the quantum mechanical side, the advantage of taking into account symmetries is well known and consists in the reduction of the full problem into a set of decoupled ones. The reduced problem has clear physical properties associated to the conserved quantities related to the symmetries because of Noether's theorem. The fact that the conserved quantities are commuting with the Hamiltonian allows for the convenient choice of the eigenbasis labeled by the respective quantum numbers that makes the Hamiltonian block diagonal. Indeed, once the selection rules introduced by the symmetries are applied, a large number of matrix elements becomes zero. From the numerical point of view, it is natural to imagine that the rewriting of the tensors in terms of the symmetry sectors allows for a more efficient computation both speed and memory-wise. Generally speaking, employing symmetries in the MPS formalism implies dealing with smaller configurational space, therefore improving the accuracy. Moreover, the targeting of specific sectors restricts the variational search to those states which have defined quantum numbers instead of approaching the problem from a "gran canonical" perspective (e.g. with the introduction of the chemical potential in the model). Clearly the latter implies a reduction of the computational complexity and results in a more precise approach than the one accounting for multiple sectors at once.

The encoding of symmetries in tensor networks is not always convenient both from a matter of efficiency and complexity though. The ones that are more straightforwardly implemented are those which transform independently each

degree of freedom of the system without making them interact. Specifically, if $U(g)$ is the unitary representation of the group element g belonging to a given group \mathcal{G} , for the latter symmetries it is possible to decompose the representation in disjointed transformations which act only locally $U(g) = \otimes_j V_j(g) \quad \forall g \in \mathcal{G}$ at site j . Concerning the complexity, a general procedure for the implementation of Abelian symmetries is explained in great detail in [37] while for non-Abelian groups the situation is quite more complicated due to the non-trivial Clebsch-Gordan coefficients.

Just to give the flavour about the implementation in the Abelian case taken from [37]: the first step is – once known the one-dimensional unitary irreducible representations in form of phase factors – to transform every link index i_r of the a tensor with $r = 1, \dots, n$ links with a tuple $i_r \equiv (l_r, \delta_r)$ where l labels the principal quantum number and δ accounts for the degeneracy. Each link then corresponds to an irreducible representation of the group labeled by l and δ . In order to make the tensor itself symmetric, the invariance property under the application of symmetry transformations to the links must be ensured. The links become directed (incoming and outgoing) to select the invariance under the direct representations and inverse (hermitian conjugate) ones. Eventually, the invariance is translated through the quantum number fusion rules to identify only some allowed matchings among the total possible combinations of the quantum numbers. Only those matches have nonzero elements in the tensor, hence the number of free parameters is reduced.

The book keeping overhead of the implementation pays out immediately when treating a system responding to the same symmetries. Indeed, all the sectors are dynamically decoupled and one can always deal with the separate, usually low dimensional, sectors without losing generality. Anyway, one must be careful when forgetting about the ensemble of the sectors as a whole for instance when the SVD is performed. Even though the symmetric matrix structure is block diagonal, the truncations during the decimation procedure should compare all the singular values and not only the ones at the sector level. Otherwise, the truncation itself could lead to artificial spurious symmetry breakings. Another crucial point is that symmetric tensors introduce constraints in the representation of a state, therefore it can happen that for some symmetry broken systems (e.g. for the case of the Néel state) the simulation can be counterintuitively more expensive.

2.8 Entanglement and correlations

For an MPS in mixed canonical form where a bond matrix λ_{α_l} is connected to two orthonormal sets of tensors as in Eq. (2.9), the bipartite entanglement properties are naturally extrapolated from the reduced density matrix. In particular, the system is then divided in left and right portions with a cut at site l of Eq. (2.9). Assuming that the number of allowed states is bounded by χ in the virtual legs, the entanglement entropy is maximized in the situation where the distribution of the coefficients is flat $\lambda_{\alpha_l} = 1/\sqrt{\chi} \forall \alpha_l = 1, \dots, \chi$. In general, the entropy of the MPS is then bounded:

$$S = - \sum_{\alpha_l=1}^{\chi} \lambda_{\alpha_l}^2 \log \lambda_{\alpha_l}^2 \leq 2 \log \chi, \quad (2.21)$$

which clearly does not scale with the size L of the chain, therefore the MPSs can only efficiently represent area-law states by construction. This also means that for the simulations of critical ground states with logarithmic corrections to the constant entropy $S \propto \text{const.} + \log(L)$, the bond dimension must be increased linearly in the system size to properly capture the long correlations. Anyway, in such a case the number of parameters is “only” polynomial in L .

Once again, the eigenspectrum of the reduce density matrix is equivalent to the squared singular value spectrum λ_{α_l} obtained from the SVD decomposition. Hence, the decimation on the SVD spectrum assumes equivalent meaning to the one in the DMRG spirit.


2.9 Boundary conditions

Boundary conditions are a crucial aspect of the simulation of lattice models and fortunately they can be handled also in the MPS formalism; indeed, for the systems investigated in this Thesis we often resort to periodic boundary conditions (PBC). There are several possibilities for the realization based on different translations of the lattice topology to the tensor network. Some tensor networks are preferable than others and in general the design has an impact on the numerical stability of the simulation. For the PBC, the most known implementation is to link the first and the last tensor of the MPS, thus taking

the trace of the tensor product:

$$\begin{aligned}
 |\psi\rangle &= \sum_{i_1, \dots, i_N} \sum_{\alpha_1, \dots, \alpha_N} A_{\alpha_N, \alpha_1}^{i_1} A_{\alpha_1, \alpha_2}^{i_2} \dots A_{\alpha_{N-2}, \alpha_{N-1}}^{i_{N-1}} A_{\alpha_{N-1}, \alpha_N}^{i_N} |i_1, \dots, i_N\rangle \\
 &= \sum_{i_1, \dots, i_N} \text{Tr}[A^{i_1} A^{i_2} \dots A^{i_{N-1}} A^{i_N}] |i_1, \dots, i_N\rangle.
 \end{aligned} \tag{2.22}$$

The tensor diagram is the following:



Unfortunately, despite ensuring a true translational invariant and periodic ansatz, this solution is not always numerically stable and neither efficient because of the loopy geometry [45] that impairs the existence of a canonical form. Only loop-free ansätze can exploit the canonical form as a huge boost to the efficiency. Another option is to consider the same trick on the Hamiltonian MPO level, the inconvenience is that during the sweeping procedure when the physical information is transmitted through the tensors' optimization, the connection between the first and the last part of the chain relies only upon the trace link. Some issues, in particular concerning the homogenization, make this realization not so effective. In order to circumvent the problem, one can employ Tree Tensor Networks [37], or – as we do here for nearest neighbour interacting models – by mapping the ring onto a chain with tailored next-nearest neighbor couplings and boundary nearest neighbor ones [46]. The loop disappears from the MPS form and gets adsorbed inside the MPO. The basic idea of such “snake” pattern is shown in Fig. 2.13 for a very little system with an even number of sites. Despite a slightly bigger MPO and a bit of extra book keeping for the post-processing of observable measurements, the numerical calculations are demonstrated to be efficient and stable [46]. Also in our work, we do not encounter any problem related to the small asymmetry of the first and last link of the chain with respect to all the others.

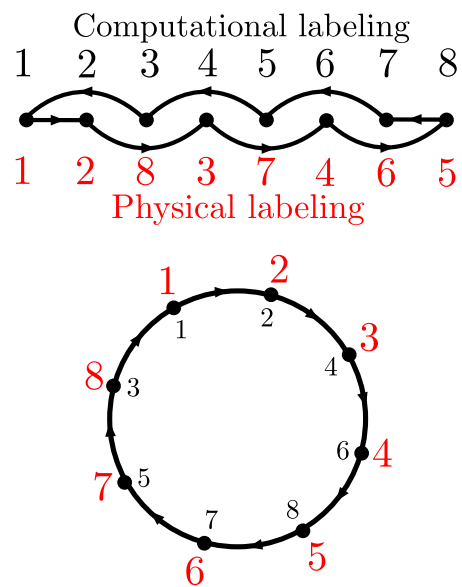


Figure 2.13: Sketch of the scheme used to mimic periodic boundary conditions with a computational open boundary conditions ansatz.

Chapter 3

Deep Learning

Deep learning is a subfield of machine learning that has been gaining a lot of attention in the last years. It comprises a set of algorithms inspired by the structures and functions of the human brain and based on so called artificial neural networks (NN). Such networks are employed nowadays in a huge variety of contexts like computer vision [47], speech recognition [48], natural language processing [49], generation [50], decision making [51] and so on, all under the hat of the artificial intelligence (AI) field (see the great textbooks [52, 53] and references therein). The success of deep learning resides precisely in the meaning of “deep”: it refers to the use of multiple layers in the NN where increasing the number of layers can be interpreted, at this point, just as increasing its complexity. The analogy with the brain is instrumental to understand the basic ideas behind the operation of a NN: its constituents are called neurons, they are connected between each other and organized in layers. The neurons in the first layer (the input layer) are connected to the neurons in the first hidden layer, the neurons in the first hidden layer are connected to the neurons in the second hidden layer and so on. The last layer contains the output neurons. When some sort of information is given to the input layer it is transmitted through the network and somehow manipulated toward the output. The incredible power of NN arises from their ability of processing the input in a non-linear way: in a simplified picture every neuron is activated when the synthesis of its inputs overcomes a certain threshold. Concretely, the neuron provides to the following ones a non-linear function (activation function) of a weighted sum of its inputs. The activation function represents the key ingredient which makes NNs universal approximators (see 4.7 of [54]). In addition to their representation power, another important advantage of

NNs is the computational efficiency. This efficiency does not refer only to the processing of the input (which is indeed just mainly a matrix multiplication algorithm) but to the optimization of the parameters of the network happening during learning. This aspect will be clarified in detail later on.

While NN framework is very general and mathematically solid, the choice of the structure of the network (or architecture) is crucially connected with the task to be solved. This is due to the fact that an AI system needs to acquire its own experience in order to “learn” how to solve a task and its way of doing it is through extraction of patterns from training data. All the learning process is nothing but a tuning (or optimization) of the model such that it is able to process the data in a way that the extracted features are the most significant ones for the problem at hand. Since the problems can be very different, the important patterns could be of completely different nature such are objects in an image or letters in a word or even more abstract concepts as the emotions contained in a speech. Therefore, the architectures able to optimally process such kind of data are different.

The last missing concept about deep learning that we want to mention in this brief introduction is the data-driven paradigm. The whole procedure for the training is aimed at making the NN build complex concepts and new representation of the data in an automated way. In order to do so, the model must analyze and interpret the provided data so it can learn from them. The outcome depends on the way in which they are processed for the sake of acquiring knowledge and being able to generalize on unseen examples.

3.1 Structure and training of a NN: concepts and a concrete example

The atom of a NN is the artificial neuron. Suppose that we have an input vector \vec{x} , the artificial neuron performs a linear combination of the input components through some weights \vec{w} (plus an eventual bias b) and applies a non-linear activation function g to the result:

$$y = g(\vec{w} \cdot \vec{x} + b). \quad (3.1)$$

The weights \vec{w} and bias b are free parameters of the neuron and since they are the target to be optimized during the training, they are called *learnable parameters*. In other words, the aim when training a model is to find the best set of parameters which maximize the accuracy of the NN prediction on the task at hand. Since the usual optimization is gradient-based, the function g must be differentiable. The most common activation functions are the sigmoid function, the softmax function, the hyperbolic tangent and the rectified linear unit (ReLU). In a NN all the neuron units are organized in a layered structure where neurons from one layer send outputs to the following one. One common instance is the *fully connected* network (FCNN) where every neuron of a layer is connected to all the neurons of the previous and following one. We denote with $y_i^{(j)}$ the output of the i -th neuron of the j -th layer and in such an architecture it can be written with the recursive relation:

$$y_i^{(j)} = g \left(\sum_k w_{ik}^{(j)} y_k^{(j-1)} + b_i^{(j)} \right). \quad (3.2)$$




Figure 3.1: Two examples of images from the training set of the MNIST dataset.

To understand the training mechanism we resort now to a concrete problem, namely the classification of handwritten digits from the MNIST database. The latter is very often taken into account to benchmark the performances of ML models for the sake of classification and consists in few tens of thousands of fixed-size images of 28×28 grey-level pixels (see Fig. 3.1 for two examples). Every image is a drawn digit between 0 and 9 and has a corresponding label. The dataset, as for all the datasets in ML, is divided into a *training set* containing those images which are used for the parameter optimization of the model and a *test set* used once the model is already trained and containing new examples which the model has not seen at training-time. The test set is of fundamental

importance to check the *generalization* ability of the trained model that is its capacity of accurately carry out the task on unseen examples. Generalization is often monitored also at training time by means of the *validation set* – containing new images with respect to the training one as well – over which the learnable parameters’ optimization is temporary frozen. More concretely, we choose as an example a training set of 60,000 images and a validation set (used also as test) of 10,000 images. We aim at carrying out a classification task: each image is given as input to a FCNN in form of a flattened vector with 784 components having values from 0 (black) to 1 (white). We design the NN so that it has as output 10 neurons whose values are in $[0, 1]$ coinciding with the probabilities of the input to belong to a given class, namely the predicted digit label. In this regard, the trained NN is supposed to emulate the probability distribution function of the classes’ list.

The explicit example of the MNIST classification problem, offers the chance to introduce some missing crucial concepts about deep learning. In order to make the NN learn to solve the classification task from the data, we want to tune the free parameters so that once one of the examples from the training set is fed to the NN, the output set of probabilities resembles the true label with the best possible accuracy. The idea behind the optimization is the minimization of a user-given metric for the accuracy of the prediction on the training data, namely the minimization of the *loss*. For the moment, we consider one natural loss function between two (discrete) probability distributions, namely the Cross-Entropy. Given two discrete probability distributions p and q on the same support \mathcal{X} , the cross-entropy of q relative to p reads:

$$H(p, q) = -\mathbb{E}_p[\log q] = -\sum_{i \in \mathcal{I}} p(i) \log q(i) \quad (3.3)$$

where \mathbb{E}_p is the expected value with respect to p . In our case we deal with a multiclass classification problem, so the support contains the digits $i = 0, \dots, 9$ (classes). Supposing that we have a certain number N of realizations $\vec{x}_{n=1, \dots, N}$, we aim at minimizing the average loss:

$$\mathcal{L} = -\frac{1}{N} \sum_{n=1}^N \sum_{i \in \mathcal{I}} p(i|\vec{x}_n) \log q(i|\vec{x}_n). \quad (3.4)$$

where we indicate with $p(i|\vec{x}_n)$ the set of probabilities of the classes for the sample \vec{x}_n in training set corresponding to the provided label. Specifically,

$p(i|\vec{x}_n)$ is encoded as a one-hot vector (a vector with all zero components but from one which is equal to one) with one on the component correspondent to the class of the label. $q(i|\vec{x}_n)$ is the output of the NN and depends on the learnable parameters that we indicate as θ standing for the collection of all the weights and the biases of the neurons.

At this point, the training can be formalized as an optimization problem to find the optimal parameters θ^* :

$$\theta^* = \operatorname{argmin}_{\theta} \mathcal{L}(\theta). \quad (3.5)$$

The high-dimensional minimization can be tackled with gradient descent when the chosen loss function is differentiable in the network parameters. Therefore, the parameters are updated with an iterative procedure by following the direction of the negative gradient of the loss. The magnitude of the shift is set with the *learning rate* α :

$$\theta' = \theta - \alpha \nabla_{\theta} \mathcal{L}. \quad (3.6)$$

Assuming a convex loss function, the gradient descent is guaranteed to converge to its global minimum. However, the problem of local minima for non-convex functions can be usually dealt with by means of the stochastic version of gradient descent. In deep learning the typical trick consists in the evaluation of the gradient and the subsequent optimization on several batches of samples stochastically drawn from the training set. For our training, we always made use of Adam [55] optimizer that is considered among the state-of-the-art techniques. In addition to the stochasticity with the batched training, it adapts dynamically the learning rate based on the first and second moments of the gradient.

Another key point of the NN training concerns the numerical efficiency. When dealing with very deep networks with a huge number of parameters, it is fundamental to perform an efficient optimization to cut the training time. In this respect, the *backpropagation algorithm* [53] is always employed and already built in the common libraries for the NNs implementation. The idea is quite easy: the dependence of the loss on the parameters is inherited from the output value of the network in the first place. The output can be written as a function of the neurons of just the previous layer and so on and so forth. All the derivatives on the loss with respect to the parameters can be decomposed following the chain rule and going backward in the networks' structure up to

the input. For simplicity, we consider the vectorized form of Eq. (3.2):

$$y^{(j)} = g(W^{(j)}y^{(j-1)} + b^{(j)}), \quad (3.7)$$

where now $W^{(j)}$ is a matrix of weights for layer j whose rows run over the neurons of the j -th layer and columns of the $j - 1$ -th one. Following the chain rule of the derivatives with respect to a generic parameter θ , it is easy to derive the following recursive equation which relates the derivative of the output $y^{(j)}$ to the ones of the previous layer $y^{(j-1)}$:

$$\frac{\partial y^{(j)}}{\partial \theta} = \frac{\partial g}{\partial \theta} \left[\frac{\partial W^{(j)}}{\partial \theta} y^{(j-1)} + W^{(j)} \frac{\partial y^{(j-1)}}{\partial \theta} + \frac{\partial b^{(j)}}{\partial \theta} \right]. \quad (3.8)$$

The essence of the backpropagation procedure is rooted in the nested structure of the derivatives. The latter can be exploited from the last layers to the initial ones by keeping track of the various contributions to the error in the prediction.




Figure 3.2: Results of the training of a fully connected network with 80 neurons in the hidden layer. The training loss is shown in red solid line, the validation loss in orange dashed line. The shaded area indicates the standard deviation of the loss over the batches of data within a given epoch. The accuracy scale is on the right y -axis and the performance is plotted in blue. The final accuracy value, namely the number of correctly classified images in the test set, is of 97.2%.

Now, with all the ingredients presented above about the learning mechanism of a NN, we explore its details on the concrete example of the MNIST classification problem. The first implemented architecture is a fully connected network which takes as input the handwritten digits images as a flattened vector of $28 \times 28 = 784$ components. The first layer has 80 neurons with a rectified linear unit (ReLU) activation function. The second and last layer consists of

10 neurons containing the predicted output probabilities for the input image of belonging to one of the classes. The network is trained on the entire training set and the weights are updated via the Adam optimizer with an initial learning rate of $\alpha = 3 \times 10^{-3}$. Each epoch contains several optimization steps, one for each batch of data in the training set (each batch was chosen to contain 100 images). At every step, the gradient of a cross-entropy loss function between the predictions of the network and the target labels are computed. In Fig. 3.2 the behaviour of the training loss, the test loss (used as validation set) and the accuracy (scale on the right y -axis) for the classification of the test set is reported for this first network. After just 7 epochs, the network is already able to correctly classify the 97.2% of the unseen data in the test set. The mean losses are represented with solid lines while their standard deviations with shaded regions around the mean.

In order to visualize what happens to the learnable parameters at training time, we report in Fig. 3.3 the evolution of the weights in the last layer of the above-mentioned FCNN (left) and the correspondent gradients. The heatmaps are histograms of the parameters (gradients) counting their number (color dimension) for a certain value (y -axis) as the training proceeds (epochs on the x -axis). The weights evolve from the random initialization and tend to stabilize in the last epoch when the majority of their gradients is zero. Meanwhile, the training loss in Fig. 3.2 decays as well as the test loss, reaching very similar values near 0.1. A lower loss in the prevision is evidently followed by a better accuracy in the classification.

The trend of the curves in Fig. 3.2 is quite optimal apart from the fine-tuning of the hyperparameters (such as the learning rate) or some subtleties concerning the choice of the non-linearities which can allow for a slightly better performance of the network. Now we present some other configurations that are instead more problematic but instrumental to see in practise how delicate can be the choice of a model. The first one leads to *underfitting*: when the number of the degrees of freedom of the network – or its complexity – is not big enough, the result of the predictions cannot reach a good level. This is represented in the left panel of Fig. 3.4 for a network with only 20 neurons in the first layer. The performance of such a network is compared with the baseline given by the previous presented model (gray and black lines). The losses and the accuracy of the 20 neurons network are visibly worse than the previous case and one could prove that even proceeding further with the training does




Figure 3.3: Evolution of parameters (left) and their gradients (right) of the last layer of the FCNN with 80 hidden neurons for the MNIST classification. The heatmaps are histograms of the parameters/gradients counting their number (color dimension) for a certain value (y -axis) as the training proceeds (epochs on the x -axis). The weights evolve from the random initialization and tend to stabilize in the last epoch when the majority of their gradients are zero. Meanwhile, the training loss in Fig. 3.2 decays as well as the test loss, reaching very similar values near 0.1.

not substantially increase the accuracy.

The other fundamental problem is *overfitting*: in this case the number of degrees of freedom of the network is too big. What usually happens in this scenario is that the training loss reaches very low values, meaning that the network really is learning to perform the wanted task (regression, classification, reconstruction, etc.) on the training set very well. On the contrary, the generalization ability of the network is more and more compromised, thus the results on the test set become worse or stop to improve. The resulting model has not built abstract concepts which allow for a good generalization on unseen data but has rather optimized its parameters in a way that is best suited for the training data only, even learning their statistical noise. The latter situation is represented in Fig. 3.4 on the right for a network with 784 neurons (the same as the input dimension) in the first layer. After the initial 1-2 epochs, the training loss keeps decreasing but this is not followed by the test loss. Indeed the accuracy becomes stagnant or even decrease. The zero-level solution to this very frequent problem is called *early stopping* and consists in stopping the training when the validation loss reaches its minimum. Already at the 2nd epoch the results are indeed better than the baseline because the model is way more powerful than the 80 neurons version. While early stopping can be a good solution to overfitting, it is not really systematic when the loss

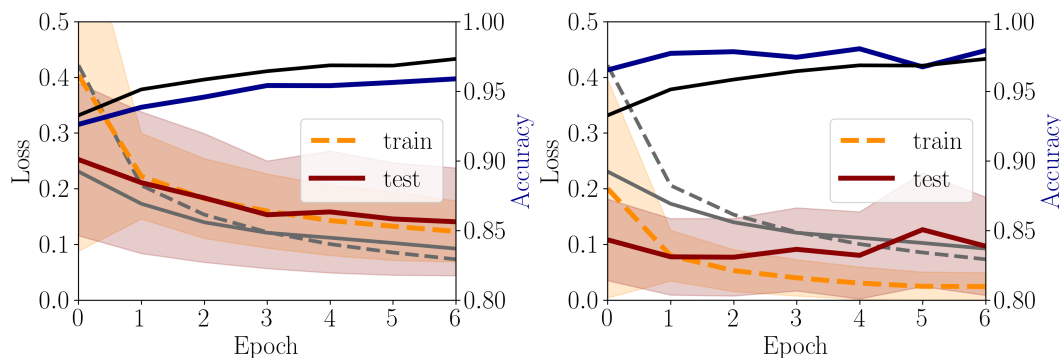


Figure 3.4: Left: underfitting. The network with 20 neurons in the first layer is compared with the baseline (gray and black lines) of the FCNN with 80 hidden neurons. The losses and the accuracy are visibly worse than the baseline and one could prove that even going ahead with the training does not substantially increase the accuracy. Right: overfitting. The network with 784 neurons in the first layer is compared with the baseline (gray and black lines). After the initial 1-2 epochs, the training loss keeps decreasing but this is not followed by the validation loss. Indeed the accuracy becomes stagnant or even decrease.

profile as the training proceeds is not as smooth and regular as in the easy cases presented here. To tackle the problem even with very complicated networks and highly dimensional loss landscapes, two other solutions are widely employed in machine learning. One is regularization: a penalty for weights with big magnitude is added to the loss therefore the network is encouraged to perform good predictions on the training set with a lower number of weights or with smaller ones. Regularization actually can be extended also for other specific constraints and provides a very useful method to gain more control on the automatic optimization of the parameters. At the same time, one must be aware that adding a lot of regularization can lead to a very biased learning where the network is too constrained. The typical regularizations for struggling against overfitting are L1 and L2 regularization [56] corresponding to penalties proportional to the absolute value of the parameters $\propto \sum_i^{N_{param}} |\theta_i|$ or to the squared parameters $\propto \sum_i^{N_{param}} \theta_i^2$. Sometimes the biases are not regularized and the above terms only contain the weights. The other clever way of systematically avoid overfitting is *dropout* [57]. The idea is quite simple: during every training epoch a random fraction of the neurons of the network is “dropped out” meaning that the neurons’ forward and backward connections are temporarily removed. By dynamically randomly changing the switched-off neurons, the network is proven to become more robust against statistical noise in the training data and to perform a better feature extraction for the prediction. The only parameter to be tuned in this case is the percentage of the




Figure 3.5: Comparison of the validation performances of the overfitting FCNN with 784 neurons (yellow) with the same network equipped with different regularization terms and dropout as in the legend. The baseline (gray) corresponds to the FCNN with 80 hidden layers. Clearly, the regularization terms allow for a better generalization of the network.

dropped neurons (up to maximum 30 – 50% of them). We report in Fig. 3.5 the comparison of the validation performances of the FCNN with 784 neurons when the regularization terms and dropout are added in order to increase the generalization ability of the network.

3.2 Autoencoders

In this section, we describe the main ingredient that we use for implementing the recognition of phases for quantum many-body systems with a machine-driven protocol, namely the deep Autoencoder (AE) architecture. The AE is a network devoted to the efficient non-linear encoding of unlabelled data, hence in the framework of unsupervised learning where the training is not carried out by adapting the prediction of the network on the “right answer” which tags the data. The main working principle of the AE is to learn a mapping from the input space to a lower-dimensional *latent space*, where the data is encoded in some automatically generated meaningful representation. The encoding is validated so that from the latent space, it is possible to reverse the transformation and faithfully reconstruct the input through another portion of the network. Being a very powerful tool, it has many applications such as dimensionality reduction, data compression, feature extraction and de-noising [53].

To address the encoding problem, AE resorts to non-linear mapping. The structure is divided in two parts: the first is called *Encoder* and takes as input a d -dimensional vector and maps it to the latent space (also called bottleneck,

of dimension k). The second part, a.k.a. the *Decoder*, re-maps the latent space back to the d -dimensional original space. The problem of creating an almost lossless compression is actually really un-trivial and the AE takes advantage of the non-linearity and the automatic way of learning from the dataset the abstract concepts in the bottleneck over other more simple methods like Principal Component Analysis (PCA)[58].

From a probabilistic point of view: consider some probability distribution $p(\vec{x})$ that has a support of dimension d . The goal of the AE is to learn an ϕ parametrized encoding function to the latent space $f_\phi(\vec{x}) = z$ and a θ parametrized decoding one $g_\theta(z) = \vec{x}'$ such that the expected value on p of the reconstruction error is minimized. We denote with $d(\vec{x}, \vec{x}')$ as the measure of the reconstruction quality. The optimal parameters ϕ and θ for the AE are found by solving the minimization problem:

$$\operatorname{argmin}_{\phi, \theta} \mathbb{E}_{p(\vec{x})} [d(\vec{x}, g_\theta(f_\phi(\vec{x})))]. \quad (3.9)$$

In practice, we do not know $p(\vec{x})$ but rather we have N samples that we assume to be drawn independently and identically distributed from it, therefore the goal becomes to minimize the empirical reconstruction error, a.k.a. the reconstruction loss. The measure of the distance $d(\vec{x}, \vec{x}')$ is usually chosen to be the squared Euclidean distance, but other choices are possible according to the nature of the input data. We will introduce afterward another appropriate candidate for dealing with probability distributions as input, which is the Kullback-Leibler divergence measuring the difference between two probability distributions. We sketch in Fig. 3.6 one possible structure of an AE made of fully connected layers.

For the sake of discerning between quantum phases, we enhanced the performances of the AE architecture with the *adversarial training*. The latter step can be implemented by adding to the AE a *discriminator* network that is trained to distinguish between the original data and the reconstructed one. The full network takes the name of Generative Adversarial Network (GAN) consisting in a generator (the AE in our case) and a discriminator. Usually, GANs are designed to generate synthetic data that closely resembles the training data they were provided with. The generator network's purpose is then to generate new data samples by learning the underlying patterns and distribution of the training data. Often it takes low-dimensional random input, referred to as noise or latent variables, and transforms it into synthetic data




Figure 3.6: Schematic representation of a fully connected Autoencoder. The Encoder f_ϕ maps the input x to the latent space z and the Decoder g_θ re-maps z back to the original space for the reconstructed version of the input x' . The training is performed by minimizing the distance $d(x, x')$ between the input and its reconstruction.

samples through a decoder. Its goal is to generate realistic data samples that are indistinguishable from the real data.

On the other hand, the discriminator network is responsible for distinguishing between real and generated data samples. It is trained on both real and synthetic data and learns to classify them correctly. During training, the generator and discriminator networks are trained simultaneously in an adversarial manner: the generator aims to fool the discriminator by generating realistic data samples that the discriminator cannot differentiate from real ones. Conversely, the discriminator aims to become more accurate in distinguishing between real and generated data samples. This adversarial process drives both networks to improve their performance iteratively.

Through this adversarial competition and iterative training process, GANs can learn to generate high-quality synthetic data samples that capture the underlying distribution of the training data. They have been successfully used in various domains, such as image generation, text generation, and even music and video synthesis. They have contributed to advancements in areas like computer vision, natural language processing, and creative applications (see the review [50] and references therein).

As mentioned above, our generator does not take noise as input but is rather a full AE. The rationale of this choice resides in the anomaly detection scheme

which we are going to describe in detail, together with the description of our GAN implementation, in the results chapter.

Chapter 4

Introduction to quantum phase transitions

This chapter is devoted to the review of some fundamental concepts about quantum phase transitions. The majority of the information presented here is derived from the textbooks by Subir Sachdev [59] and by Amit Dutta et al. [60]. Rather than conducting a thorough exploration of the properties of quantum phase transitions, we aim to provide a general overview that emphasizes some keywords that will be encountered later in the thesis.

In our everyday life, we typically observe classical phase transitions, such as for instance the water-vapor transition. Along such transitions, the physical properties of a given medium undergo a transformation in response to alterations in external conditions controlled by a parameter like temperature or pressure. This transformation can be abrupt and marked by a discontinuous change in some thermodynamic quantity, referred to as first-order phase transitions, or it can be continuous. The identification of the external conditions at which the transformation occurs defines the phase transition point or critical point. While discontinuous changes are generally associated to absorption or release of latent heat, in the continuous case, due to thermal fluctuations, the phases becomes indistinct at the critical point (critical opalescence). Following the example of [60] we can consider a translational invariant system in d dimension with a critical temperature T_C . The transition from one phase to another is usually also connected to an order parameter $O(x)$ that has a thermal expectation value of zero below T_C and is finite in the other phase. Near the critical point, the order parameter vanishes with a power-law dependence on

the distance from the critical temperature $\langle O(x) \rangle \sim |T - T_C|^\beta$. Also the two-point correlation function of the order parameter at different spatial positions $\langle O(x_1)O(x_2) \rangle$ that has usually an exponential decay at large distances in the high temperature phase, changes its behaviour at the critical point as a consequence of the divergence of its typical length scale set by the correlation length ξ . The latter, indeed, diverges with $\xi \sim |T - T_C|^{-\nu}$, where ν is another critical exponent. As a consequence, at the critical point the concept of typical scale is lost and the system becomes scale invariant. Furthermore, other thermodynamical quantities such as the specific heat and the susceptibility diverge at the critical point with algebraic scalings whose exponents are related to the ones above. The connected power-law dependencies already represent a quite notable feature, but there is more! Perhaps, the most striking property of phase transitions is that even the numeric values of the critical exponents depend only upon the dimensionality and the symmetries of the system and of the order parameter, allowing for a classification based on classes of equivalence also named as universality classes. The critical behaviour of all the models within a given universality class is the same and irrespective of the microscopic details of the model itself. It is worth pointing out that the abovementioned divergences and discontinuities occur only at the thermodynamic limit because all the thermodynamical quantities for a finite system are analytical functions of their variables.

Unlike thermal phase transitions, which are driven by thermal fluctuations, most of the quantum phase transitions are caused by quantum fluctuations that arise from the Heisenberg uncertainty principle. Compared to their thermal analogue, quantum phase transition can occur at zero temperature where the system is in its ground state. The control parameter becomes now a coupling constant g in the Hamiltonian that determines an imbalance between different non-commuting terms. Whereas deep in a quantum phase the physics is dominated by a single term resulting in a ground state which is also its eigenstate, in an intermediate coupling regime the physics is governed by the interplay of various states in form of a superposition of the non-commuting terms' eigenstates. This also represent another motivation for the study of such phenomena from the interacting many-body systems' perspective: the most explored regimes are usually the weak coupling Hamiltonian and the strong coupling limit so the study of the intermediate region is left apart. Moreover, the quantum nature of the fluctuations generates a more rich phenomenology in comparison to the classical case and many more universality classes have been identified.

Quantum phase transitions are associated to non-analiticity points of the ground state energy of the infinite system as a function of g , instead of the free energy of their classical counterpart. At the quantum critical point (QCP), in most cases the characteristic energy scale of the fluctuations Δ (being for example the energy gap between the ground state) vanishes as $\Delta \sim |g - g_C|^{\nu z}$ where ν and z are again critical exponents. For the finite system the gap does not really close but the transition manifests itself with an *avoided level crossing* between the ground state and the first excited state's energies. The order parameter is now a quantum expectation value and scales as $O \sim |g - g_C|^\beta$ near the QCP.

While the dynamics of a classical system over time is not a significant factor in the analysis of thermal phase transitions at equilibrium, for the quantum case, space and time are connected with each other. We should indeed make a distinction when dealing with correlation functions: on the one hand there are the equal-time connected correlation functions for a distance r and, on the other, the equal-space version for different times. The former includes the definition of the characteristic correlation length scale ξ far from the QCP:

$$G(r) = \langle O(0, t)O(r, t) \rangle - \langle O(0, t) \rangle \langle O(r, t) \rangle \propto \frac{e^{-r/\xi}}{r^{d-2+\eta}} \quad (4.1)$$

where η is the anomalous dimension of the order parameter (or Fisher exponent). The qualitative change in the nature of the ground state spatial correlations along the QCP is usually accompanied by a power-law divergence of the spatial correlation length $\xi \sim |g - g_C|^{-\nu}$. The latter has a different scaling with respect to the equal-space correlation length ξ_τ but the two are anyway related via the already encountered z factor by $\xi_\tau \sim \Delta^{-1} \propto \xi^z$. Together with the qualitative change in the correlations, also the dispersion relation of fluctuations is strongly modified at the critical point due to the closure of the gap. The closure of the gap, indeed, provides the system access to excitations that cost an infinitesimal amount of energy and therefore are easily populated.

Both for thermal and quantum, the emergence of an order parameter that is finite on the “ordered” side, is associated with a spontaneous symmetry breaking of the (equilibrium or ground) state meaning that the chosen configuration does not respect the symmetries of the underlying physical theory. It is worth mentioning that symmetry breaking is not a necessary condition for a phase transition to happen though. In fact, for classical systems with $d \leq 2$ or

one-dimensional quantum systems with sufficiently short-range interactions – particularly relevant for this thesis – the spontaneous breaking of a global continuous symmetry is not even possible as a consequence of the Mermin-Wagner theorem. In such low-dimensional systems, long wavelength fluctuations cost little energy and are able to destroy the true long-range order. However, the theorem does not prevent the existence of a phase transition in the sense of a diverging correlation length ξ such as for the Berezinskii-Kosterlitz-Thouless transition where there exist a high-temperature/disordered phase with exponential falloff of the correlation function and a low-temperature/ordered phase with power-law decay (quasi long-range order).

Part II

Results

In this part, we present our results for the data-driven approach to the detection of phase transitions in quantum many-body systems. The content is a unification of our publications [61, 62] as well as additional outcomes of our research and some results contained in [63].

Traditionally, phase transitions have been characterized using analytical methods possibly exact – in a very limited set of cases – or through mean field approaches; or more and more often resorting to numerical simulations. The detection normally consists in the measure of specific quantities (local observables, correlation functions, energy gap, etc.) which reflect the changes in the physics through the transition from one thermodynamical state to another. The main concern of such an approach is the prior knowledge about the system needed for computing the order parameter that is associated to its universality class. Very recently, another strategy has gained more and more attention in the physics community: the idea is to consider raw and very general data extracted from the system and distinguish between different patterns within the data by resorting to machine learning techniques [64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75]. As already introduced in the previous chapter, the relevance of these techniques is rooted in the ability of a properly trained machine to process extremely complex data in an efficient way for the sake of pursuing classification tasks, pattern recognition, generation of brand new data responding to some constraints and even development of decision processes. Our proposal is to follow this new promising data-centric perspective with a two-fold contribution: on the one hand, we improve the recently developed machine learning architecture for the automated unsupervised detection of phase transition by employing state-of-the-art models from the artificial intelligence community; on the other hand we propose and study a new, very general and experimentally measurable, quantity that can be used as a proxy for the scope.

The chapter is organized as follows: in the first part we present the models of the quantum systems on which we tested our detection algorithms. These were numerically simulated for obtaining the ground state by means of MPS ansätze in order to have access to the relevant quantities. The latter models are the XXZ spin chain, the extended Bose-Hubbard model and the two-component Bose Hubbard model. Then we present the Generative Adversarial Network (GAN) that we use for detecting changes in the ES patterns along a very subtle phase transition, namely the Berezinskii-Kosterlitz-Thouless (BKT) tran-

sition, through an anomaly detection scheme. The results of the detection are reported for all the three models together with an analysis of the ES in the different regimes. Afterward, we show the characterization of the candidate proxy quantity: the full probability density function (PDF) of a globally $U(1)$ conserved charge in an extensive sub-portion of the system. We focus in this case on the extended Bose Hubbard model, hence analyzing the PDF of the boson number in a sublattice, and we discuss a simple educated fit of the PDF that allows for a complete mapping of the phase diagram. In the final part, we describe another machine-learning architecture, namely the Concrete Autoencoder, that can be used to perform an automated feature selection task on the input data. We show that the model is capable of suggesting the most relevant eigenvalues of the ES in order to reconstruct the entire spectrum in one phase of the system. Therefore, this model is highly interpretable and its use can provide both a means for automatically distinguishing between different phases and an indication of which are the most important features in the data for characterizing those phases.

Chapter 5

Models

Below we list the main features of the models which we consider in this thesis for the purpose of detecting quantum phase transitions. The description is not intended to be completely exhaustive but it is rather a summary of the models' relevant features for the problem at hand. For the details we point to specific works or reviews (and their references) in the literature.

5.1 XXZ model




Figure 5.1: Sketch of the ground state phase diagram of the XXZ model. The three phases are: ferromagnet (FM), paramagnet (PM) and antiferromagnet (AFM).

The XXZ spin chain is an exactly solvable model that generalizes the Heisenberg chain by introducing an anisotropy in the spin interaction. It has been extensively studied in the literature (see the very good reviews in [34, 76]). For this reason, we use this model as a benchmark since its phase diagram is precisely known. The Hamiltonian of the XXZ model is given by:

$$H_{XXZ} = -J \sum_{j=1}^{L-1} \left[\frac{1}{2} (S_{j+1}^+ S_j^- + S_j^+ S_{j+1}^-) + \Delta S_{j+1}^z S_j^z \right], \quad (5.1)$$

where $S^\alpha = \frac{1}{2}\sigma^\alpha$ (σ^α being the Pauli matrices) and J is positive and sets the scale of the energy in our notation. At zero temperature, by changing the anisotropy Δ , the model has three different phases. For $\Delta > 1$ the ground state is gapped and exhibits ferromagnetic order where all the spins are aligned in either $+z$ or $-z$ direction; in the $\Delta \rightarrow \infty$ limit it is a product state and the low-energy excitations are magnons. This is a phase with broken symmetry because the state does not exhibit the discrete \mathbb{Z}_2 symmetry of spin reflection under which the Hamiltonian is invariant. The Heisenberg chain is recovered when $\Delta = 1$. In this limit, the \mathbb{Z}_2 symmetry is enlarged to the full rotational symmetry and, as a consequence of the Goldstone's theorem, the ground state has a gapless spectrum. For $|\Delta| < 1$, the configuration is paramagnetic with zero magnetization, the correlations show power-law behaviour thus the phase is critical. In this phase the low energy is effectively described by the Tomonaga-Luttinger liquid [77] (see also appendix A). In the $\Delta < -1$ regime, the ground state is gapped and has antiferromagnetic order. When $\Delta \rightarrow -\infty$ in the thermodynamic limit, the ground state is again symmetry broken is chosen from one of the two degenerate Néel states $|N_1\rangle = |\uparrow\downarrow\uparrow\downarrow\dots\rangle$ and $|\downarrow\uparrow\downarrow\uparrow\dots\rangle$. The total magnetization is zero in this case but the staggered one is finite. The low-energy excitations are domain walls or flipping in the Néel order such that two spins are consecutively aligned. The $\Delta = -1$ configuration is also known as isotropic Heisenberg antiferromagnetic point and is probably the most interesting regime of the XXZ model. When this point is crossed from the paramagnetic side, a BKT phase transition from the gapless paramagnet to the gapped antiferromagnetic state occurs: the gap opens up with a non-analytic dependence on the anisotropy Δ corresponding to a correlation length [76] $\xi \propto e^{-\frac{\pi}{\sqrt{\Delta+1}}}$. The phase diagram of the XXZ model is shown in the sketch of Fig. 5.1.

5.2 Extended Bose Hubbard model

The one-dimensional Extended Bose Hubbard is a non-integrable model that also has been thoroughly analyzed in literature: detailed information on both the numerical and experimental aspects about this model can be found in [78, 79, 80, 81, 82, 83, 84, 85, 86] and references therein. The Hamiltonian for an




Figure 5.2: Sketch of the Extended Bose-Hubbard ring with hopping parameters t , on-site interaction U and nearest-neighbour interaction v .

L -sites system is:

$$H = \sum_{j=1}^L \left(-t(b_{j+1}^\dagger b_j + \text{h.c.}) + \frac{U}{2} n_j(n_j - 1) + V n_j n_{j+1} \right), \quad (5.2)$$

where $b_j^{(\dagger)}$ is the annihilation (creation) bosonic operator and $n_j = b_j^\dagger b_j$ the number operator on site j along a one-dimensional chain. The hopping coefficient is denoted by t and is set to 1 giving the energy scale in our notations, while U and V are the on-site and nearest-neighbour interaction strengths. The zero-temperature phase diagram of the unitary integer filled lattice, $\nu = N/L = 1$, presents a number of prototypical quantum phase transitions [84, 85], namely (i) a BKT transition between a gapless superfluid (SF) and a gapped Mott insulator (MI) phases, (ii) a quantum phase transition between the MI and a topological Haldane insulator (HI), and (iii) an Ising type transition between the HI and a charge density-wave (CDW) state. It has been recently pointed out that a phase-separated regime between SF and supersolid (SF + SS) [73, 86] is present in the bottom right corner of the phase diagram (for large V and small U).

In the $V = 0$ case, the model resembles the well-known Bose-Hubbard model with the famous BKT transition between the superfluid and the Mott insulator. The phase transition results from the competition between the delocalization effects of the kinetic term, which reduce the phase fluctuations, and the localization effects of the interaction term, which reduce the on-site particle number fluctuations. The Mott insulator has an energy gap so that the smallest cost for a particle-hole excitation is finite. The gap opens with the typical BKT behaviour as [83]:

$$\Delta \propto e^{-\text{const}/\sqrt{(1/U_c)-1/U}} \quad (5.3)$$

that makes the determination of the critical value U_c much more difficult in the one-dimensional setup. The superfluid configuration has the features of the Tomonaga-Luttinger liquid: the low energy excitations are collective modes (phonons) with linear dispersion relation and the correlation functions show an algebraic decay with an exponent that depends on the model parameters. On the other hand, the insulating phase lacks of such an order and is characterized by an exponential suppression of the $\langle b_j^{(\dagger)} b_{j+r} \rangle$ correlator [84]. The addition of the nearest neighbour interaction term enriches the physics: in the CDW, the phase is characterized by a staggered diagonal order $(-1)^r \langle \delta n_j \delta n_{j+r} \rangle$ where δn_j is the boson number fluctuation from average filling that decays to a constant value. For very large V the typical pattern of density wave ...202020... for the occupation number appears. The extended interactions can also stabilize the HI phase which possesses hidden order that can only be revealed by non-local string correlations [84]. The presence of edge modes in this phase can also be observed in some degeneracies in the ES [85]. The nature of the excitations in this exotic phase consists in an interplay between charge excitations and neutral ones: this has been studied in [84] by observing the single particle gap, which is finite in all the insulating phases and closes at the MI-HI transition; and the neutral gap which closes both at the MI-HI and HI-CDW transitions. Finally, for the present unitary filling case, a phase-separated region consisting of a supersolid and superfluid part has been recently studied in [86], as signaled by power-law decay of $\langle b_j^{(\dagger)} b_{j+r} \rangle$ and uniform density in the superfluid part whereas the supersolid region still has quasi long-range order for the correlator but staggered local density.

5.3 Two-component Bose-Hubbard model

The last model that we study is a generalization of the Bose Hubbard model for two species. Cold atomic gases with two hyperfine levels (denoted in the following with A and B) confined in optical lattices of L sites can be described

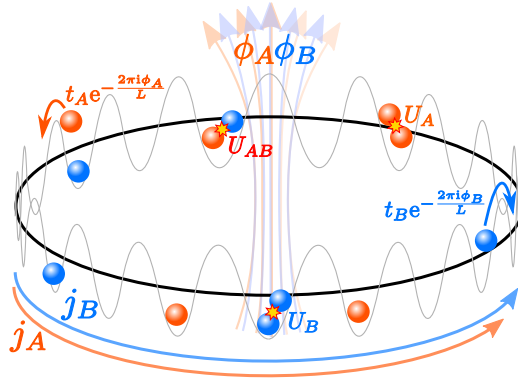


Figure 5.3: Sketch of the 2-component Bose-Hubbard ring with hopping parameters $\tilde{t}_\alpha = t_\alpha e^{-i2\pi\phi_\alpha/L}$ (twisted periodic boundary conditions), on-site intra-species interactions U_α and inter-species interaction U_{AB} . The two fluxes ϕ_α pierce the ring and give rise to bosonic currents j_α in the ground state and allow to study the superfluid properties of the system. We limit our study of the ES for the phase detection task to the $t_\alpha = t$ a real number.

by the following Hamiltonian:

$$\begin{aligned}
 H &= H_A + H_B + H_{AB}, \\
 H_\alpha &= \sum_{j=1}^L \left[- \left(t_\alpha b_{j+1,\alpha}^\dagger b_{j,\alpha} + \text{h.c.} \right) + \frac{U_\alpha}{2} n_{j,\alpha} (n_{j,\alpha} - 1) \right], \\
 H_{AB} &= U_{AB} \sum_{j=1}^L n_{j,A} n_{j,B},
 \end{aligned} \tag{5.4}$$

where the operators $b_{j,\alpha}^\dagger, b_{j,\alpha}, n_{j,\alpha}$ are the analogue of the Bose-Hubbard case but they are labeled with the flavour index $\alpha \in \{A, B\}$. The single species Hamiltonian H_α accounts for the hopping between neighboring sites and for the on-site repulsion characterised by the parameter $U_\alpha > 0$, therefore matching Eq. (5.2). The Hamiltonian H_{AB} describes the inter-species on-site interaction with strength U_{AB} . We limit ourselves to a zero temperature, \mathbb{Z}_2 symmetric mixture where $t_\alpha = t$, $U_\alpha = U = 10t$ and filling $\nu_\alpha \equiv N_\alpha/L = \nu/2$, in terms of the number of atoms N_α . The phase diagram of the 1D model is very rich and has not yet been determined with the same accuracy as in higher dimensions [87, 88]: to our knowledge the most complete analysis can be found in [89]. For the purpose of the present thesis we focus on a small region of the parameter space where only two phases appear: the two-superfluid (2SF) phase and the pair-superfluid (PSF) phase. The 2SF phase is characterised by both components A and B being superfluid. The low energy spectrum consists

of two gapless linear (Goldstone) modes corresponding to a density (in-phase) and a spin (out-of-phase) mode. In the PSF phase, the two components are paired and the spin-channel acquires a gap. This phase is sometimes considered the bosonic counterpart of the much more relevant phenomenon of pair condensation occurring in fermionic systems.

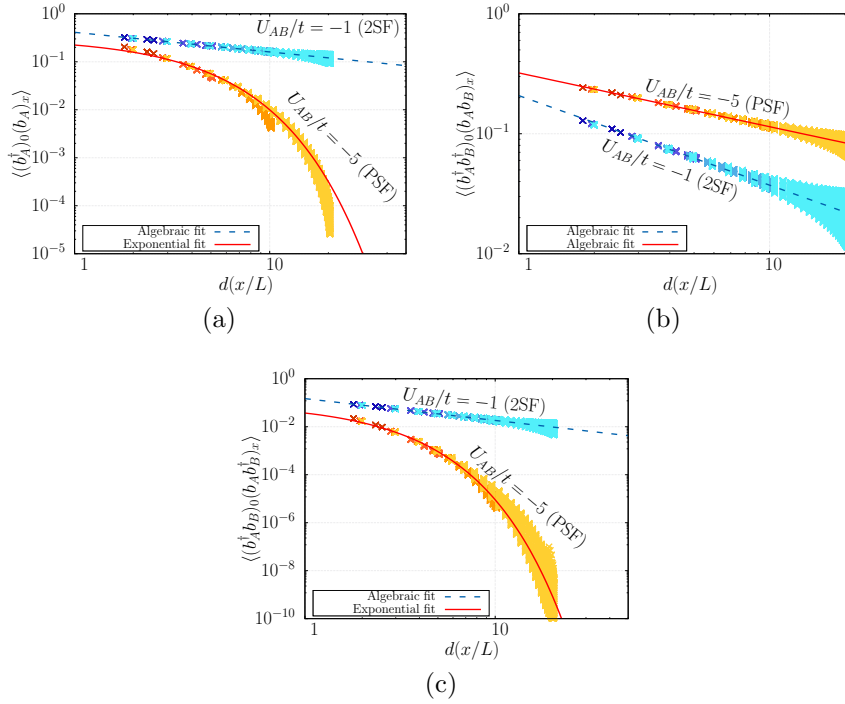


Figure 5.4: The correlation functions of Eq. (5.5) (the (a) panel is G_α , (b) is R_D and (c) is R_S) as function of the conformal distance for a unitary total filling system $\nu = 1$, $t = 1$ and $U = 10$. The orange points concern a regime in which the system is in PSF phase while the blue ones concern the 2SF phase and we represent with a color gradient from dark to light different system's sizes from $L = 8$ to $L = 64$. The dashed lines are exponential and algebraic fits depending on the expected behaviour of the functions for $U_{AB}/U = -0.1$ and the solid one for $U_{AB}/U = -0.5$.

In order to study the superfluid properties of the system we made use in [63] of a standard technique involving twisted periodic boundary conditions [90]. The hopping between neighbouring sites in the single species Hamiltonian is modified as $\tilde{t}_\alpha = t_\alpha e^{-i2\pi\phi_\alpha/L}$ and $t_\alpha, \phi_\alpha \in \mathbb{R}^+$ so that an artificial flux ϕ_α pierces the periodic one-dimensional chain (a ring). In [63] we have shown that the addition of the latter term results in persistent superfluid currents whose behaviour is affected by the inter-species interaction and related to the so called Andreev-Bashkin effect [91] that is not going to be discussed in this Thesis though. Having access to the superfluid densities of the system

allowed us to characterize the transition between the 2SF and PSF regime. The prediction for the critical value of U_{AB} when it occurs are compared with our machine-driven analysis. In [63] we also took into account the scaling of correlation functions in those phases, in particular:

$$\begin{aligned} G_\alpha(x) &= \langle b_{i+x,\alpha}^\dagger b_{i,\alpha} \rangle, \\ R_D(x) &= \langle b_{i+x,A}^\dagger b_{i+x,B}^\dagger b_{i,B} b_{i,A} \rangle \\ R_S(x) &= \langle b_{i+x,A}^\dagger b_{i+x,B} b_{i,B}^\dagger b_{i,A} \rangle. \end{aligned} \quad (5.5)$$

The single-body correlations, G_α , have a mixed density/spin (in/out-of-phase) character. The two contributions can be instead isolated with the help of two-body correlations: R_D concerns the superfluid character of pairs of $A - B$ particles and therefore the density channel, while R_S relates to particle-hole pairs and therefore the spin channel [89, 92].

Away from commensurate effects, possibly leading to a Mott insulator, the density channel is always superfluid, i.e., R_D scales algebraically (a.k.a. quasi long-range order [77]). A change in R_S and G_α from algebraic to exponential decay happens instead when entering the PSF phase, due to the opening of a gap in the spin channel. This is illustrated in Fig. 5.4, where the correlations measured for different system sizes ($L = 8, 16, 32, 64, 96$) are reported for two sample parameter values deep in the 2SF (blue) and PSF (orange) phases. The numerical value of the exponent of the power-law decay is associated to the Luttinger parameters of the effective low-energy theory of the system (see [89] and appendix A).

5.4 Numerical details

We employed MPS ansätze with embedded symmetries (in particular $U(1)$) for computing the ground state of all the systems and extracting the ES or the PDF. The XXZ system, the simple BH without the nearest neighbour interaction and the BH2S were simulated for [63] with open boundaries for XXZ and BH and with periodic boundaries for BH2S. The EBH with the full interaction term for computing the PDF was instead considered afterwards and with periodic boundary conditions. All the periodic systems are implemented by means of the “snake” pattern in the couplings of the MPO (see previous Chapter for the details).

For the bosonic models, we set the maximum occupation to $N_{max} = 4$ bosons per site (EBH) and $N_{max}^{(tot)} = (N_A + N_B)_{max} = 6$. Setting a proper number of states through the bond dimension of the MPSs, we can reliably compute relevant quantities for several system sizes, with a discarded probability at most of 10^{-9} .

Chapter 6

Detection of the BKT transition with GAN

6.1 Anomaly detection scheme

Anomaly detection is a process of identifying unexpected patterns or behaviors in data that deviate from the normal pattern. The goal of anomaly detection is to automatically recognize instances in data that are unusual or represent some type of outlier and therefore may be indicative of a problem, fraud, or error. This is a crucial task in various domains and applications where it can help prevent financial loss, security breaches, health problems and so on (see [93] and references therein).

The basic idea behind anomaly detection is to learn the normal patterns in data and then flag any observations that fall outside the normal range as anomalies. A machine learning based anomaly detection method can be classified into one of two main categories: supervised or unsupervised. As usual, a supervised method uses labeled data to train a model that is then used to predict anomalies in new, unlabeled data. Unsupervised anomaly detection methods (the ones which we use here), on the other hand, do not use labeled data and rely on three main approaches: statistical methods, clustering, and neural networks. Statistical methods employ basic statistical measures such as mean, standard deviation, and probability distributions to identify anomalies. For example, entries in a dataset drawn from a Gaussian distribution that show values lying more than three standard deviations away from the mean

might be considered anomalies. Clustering methods group similar data points into clusters, and anomalies are identified as points that are significantly different from the other points within the clusters. Also neural networks, such as autoencoders and recurrent neural networks, can be used for solving the problem. Their ability to learn the underlying structures in the data allows to identify anomalies as points that do not fit the learned structure of the normal data.

Here we report the anomaly detection algorithm we use in [61] that is based on the idea in [73] but resorts to an improved architecture, namely a Generative Adversarial Network (GAN). Our aim is to map the phase diagram of a many-body system using the anomaly detection scheme as in the following: we train the GAN on some unlabeled, representative data of the state of the system in one phase, hence corresponding to some region in the space of the Hamiltonian’s parameters. A test-time we provide the model some data from other regions of the parameter space: different phases with respect to the training one will be anomalies for the model.

6.2 GAN architecture

The anomaly detection protocol that we implemented is based on a GAN architecture composed by a full Autoencoder (AE) and a discriminator. Before addressing the role of the discriminator, we briefly describe how the AE can be used to detect anomalies. The parameters of the encoder ϕ and decoder θ (see previous section 3.2) are optimized during the training on samples that belong to a chosen region of the phase diagram (normal samples). After the training, the AE has learnt some hidden representation of the data in the bottleneck. Our assumption is that once data from another phase are fed to the AE, their hidden representation will be different making the AE lose the ability of reconstructing them with a good accuracy. The “degree of anomaly” or anomaly score can be quantified directly from the reconstruction loss $\mathcal{L}_{rec} \equiv d(x, x')$ between the input and the output of the AE.

In the exploration of a phase diagram, this scheme has several advantages because, for instance, it allows to choose the region of the normal data where their production is numerically favourable or where the physics of the system is well understood. Nevertheless, it allows for a scan of the phase diagram in an

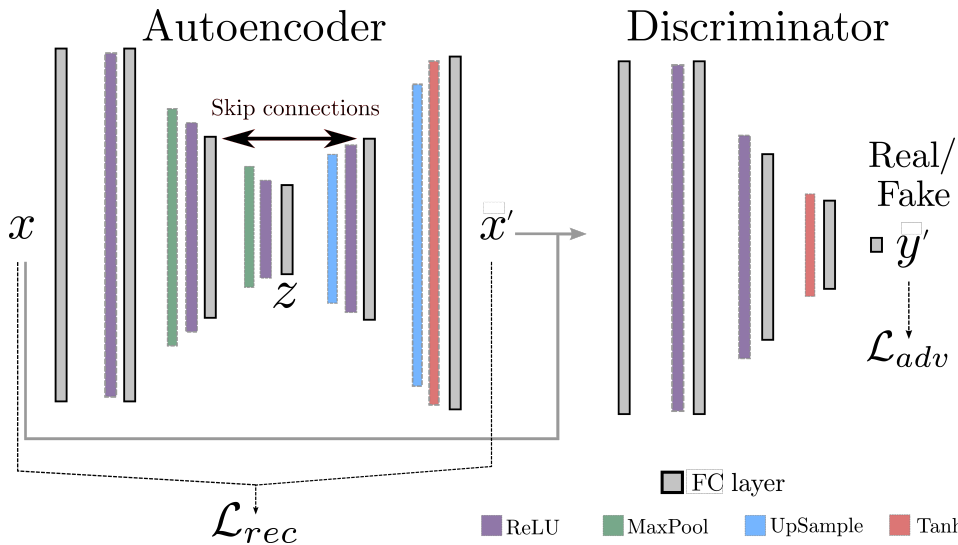


Figure 6.1: Sketch of the GAN architecture we used for the phase transition detection. The Autoencoder (AE) takes in input a vector with the eigenvalues of the ES (x) and operates an encoding and a decoding to its reconstructed version x' . The discriminator takes in input x or \hat{x} and gives a real number y' .

unsupervised way because no prior knowledge of the possible phase labelling is required.

The standard training procedure for the AE involves the minimization of the reconstruction loss but it has been demonstrated that the adversarial training usually improves the optimization of the parameters and entails a better reconstruction [53]. We decide to provide the AE with a discriminator that takes x or x' as input and outputs a single real number y' . The composite structure is a GAN [50] of which we show an example in Fig. 6.1. The role of the discriminator is to distinguish if the vector in input is a real example (one given as input) or a fake one (the reconstructed examples by the AE). For real examples the output of the discriminator should reach $y' = 1$ whereas for fake examples $y' = 0$. The two networks are trained together in a competitive way and every training step is made of two sub-steps:

- the discriminator is optimized alone in order to be able to predict that N_{batch} samples from the dataset are real and that their reconstructions by the AE are fake.
- the AE – that is a full-fledged generator of samples – is optimized in order to fool the discriminator, hence making it predict that the reconstructed samples of the batch are real.

Algorithm 1 Training of GAN

-
- 1: Initialize the parameters θ_{AE} of the autoencoder network AE
 - 2: Initialize the parameters θ_D of the discriminator network D
 - 3: Initialize the weights λ and ϵ for the losses
 - 4: **for** number of training epochs **do**
 - 5: **for** k steps **do**
 - 6: Forward the batch samples x_i , $i \in [1, N_{batch}]$ in the AE
 - 7: Update the discriminator by ascending the stochastic gradient:

$$\nabla_{\theta_D} \frac{1}{N_{batch}} \sum_{i=1}^{N_{batch}} [\log D(x_i) + \log(1 - D(x'_i))]$$

- 8: Update the AE by descending the stochastic gradient:

$$\nabla_{\theta_{AE}} \frac{1}{N_{batch}} \sum_{i=1}^{N_{batch}} [\epsilon \mathcal{L}_{rec}(x_i, x'_i) + \lambda \log(1 - D(x'_i))]$$

- 9: **end for**
 - 10: **end for**
-

The pseudo-code of the training procedure is reported in Algorithm 1. Since the discriminator can be seen as a binary classifier, the gradient ascent at line 7 of Algorithm 1 can be converted to a minimization of the following adversarial loss:

$$\mathcal{L}_{adv}(y, y') = y \log(y') + (1 - y) \log(1 - y'), \quad (6.1)$$

where the target values of the classifications are assigned as $y = 1$ for real samples and $y = 0$ for fake ones when the discriminator is trained, whereas $y = 1$ for the reconstructed samples is set when the AE is optimized. As in common practise for actualizing the adversarial competition, we separately train the sub-networks by freezing the respective parameters' optimization. Since in general the reconstruction loss and the adversarial loss range within intervals of different magnitudes, we combine them adding some weights in the overall loss (see line 8 of Algorithm 1):

$$\mathcal{L}_{tot} = \lambda \mathcal{L}_{adv}(y, y') + \epsilon \mathcal{L}_{rec}(x, x'). \quad (6.2)$$

The hyperparameters λ and ϵ are tuned in order to balance the influence of the individual losses as well as the learning rates of the two subnetworks' optimizers. At test-time the discriminator is not necessary anymore: the AE – that have benefited from the adversarial training – faces the whole phase

diagram and the anomaly score is computed. We employ a GAN architecture implemented with PyTorch [94] and the code is public at [95]. Among all the possible structures for the layers, we tried to adopt a Convolutional Neural Network (CNN) whose principle is to convolve the layers with some filters but we find that the dense structure outperforms the CNN for our task. Probably, the full-connectivity among the neurons in the former structure plays an important role for the success of the reconstruction. We also employed MaxPool filters for the down-sampling of the input in the encoding process (by considering the biggest values and neglecting the smallest ones in a fixed size sliding window) and for the decoding part the UpSample is used for going back to the original space (it fills the holes created by the up-sampling to a bigger dimension through customized interpolations). Moreover, a skip connection between two latent layers is added to improve the model performance. It consists in the transfer of an encoding intermediate layer’s information, i.e. for example x_E , to the same-size layer in the decoder x_D before the activation function. In Fig. 6.1 this is represented with an arrow. Skip connections were first introduced to prevent gradient vanishing in deep networks and then in residual networks in order to overcome the so-called degradation problem (see [54, 96] and references therein). Supposing that x_E is transformed toward the decoder with a function $F(x_E)$ and the symmetric layer in the decoder $x_D = \sigma(F(x_E))$ is obtained by applying an additional non-linearity σ , the skip connection

$$x_D = \sigma(F(x_E) + x_E)$$

does not introduce neither extra-parameters nor computational complexity. Nevertheless, the optimization of the new F (namely the *residual* mapping) is easier thanks to the information of x_E and was found to lead to the learning of good abstract representations in the bottleneck, rather than local details of the data [97]. All the manipulations on the data are made through standard built-in methods of the PyTorch library.

6.3 Features of the symmetry resolved ES

The reason why we choose the ES as a dataset is that it is a well-known quantity that exhibits distinct properties both in gapless and gapped phases [22, 98, 99], and its degeneracy pattern is often used to discriminate different kinds of

topological phases [23, 100]. Its structures embed information about non-local quantum correlations, as formalized within the bulk-boundary correspondence framework [101, 102, 103]. Both ad-hoc defined order parameters [104, 105] and machine learning driven approaches [68, 73] have been employed for the detection of phase transitions. Before showing the results of our anomaly

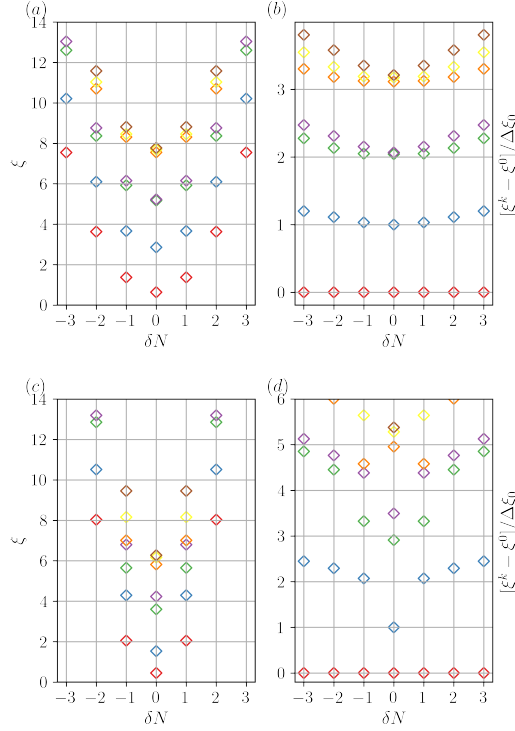


Figure 6.2: Entanglement spectrum for the $L = 600$ open XXZ chain at $\Delta/J = -0.5$ ((a) and (b)) and for the $\Delta/J = -1.2$ ((c) and (d)). In (a) the parabolic envelopes' structure is the marker of the emergent CFT. Panel (b) – (d) shows the same data of (a) – (c) obtained by subtraction of the lowest parabola $\xi^{k=0}$ and by setting the distance between $\xi_{\delta N=0}^{k=1}$ and $\xi_{\delta N=0}^{k=0}$ equal to one in order to better appreciate the equally spaced structure and the degeneracies.

detection scheme with the GAN, we describe our dataset in a quite detailed way. Specifically, we show here the characteristic features of the ES for all the three models presented above in a region of the phase diagram where the BKT transition occurs. Our analysis is focused on precise regions of the control parameters: $-1.5 \lesssim \Delta \lesssim 0$ between the paramagnetic (PM) gapless region and the anti-ferromagnetic (AFM) one for the XXZ chain, $0 \lesssim U/t \lesssim 5$ and $V = 0$ between the SF and the MI for the BH model and $-0.4 \lesssim U_{AB}/U \lesssim 0$ ($U = 10t$) for the BH2S, where the system is either in double superfluid regime or in pair superfluidity.

In Fig. 6.2 and Fig. 6.3, we plot the logarithm with base 10 of the reduced

density matrix eigenvalues $\xi = -\log_{10}(p_i)$ for a bipartition of the systems in two halves. Those are related to the eigenvalues of the entanglement Hamiltonian as presented in the previous Chapter and the base 10 is a convenient choice for the readout of the order of magnitude. The eigenvalues are labeled by the symmetry sector, corresponding for the XXZ to the number of spinless interacting fermions resulting from the mapping of the Hamiltonian under Jordan-Wigner transformations [34], the number of bosons for the BH and the number of bosons for each species for the BH2S. All the numbers are referred to the total amount of particles in one of the bipartitions. Since there are multiple eigenvalues for each sector, they can be ordered in a decreasing way: we label every eigenvalue with the (shifted) sector number, $\delta N = N - N_{mean}$, and with the sorting index, k .




Figure 6.3: Entanglement spectrum for the $L = 256$ open BH chain at $U/t = 1$ ((a) and (b)) and for the $U/t = 4$ ((c) and (d)). The parabolic structures in the gapless phase are very clean with respect to the ones of the XXZ model (see Fig. 6.2): we observed a very slow convergence to the thermodynamic limit for the spin chain indeed.

The XXZ and the BH model exhibit parabolic structures at fixed k and the parabolas with different k have the same curvature. Besides, the parabolas are translated by multiples of the difference between the first and the second

eigenvalues of the $\delta N = 0$ sector:

$$\Delta\xi_0 = \xi_{\delta N=0}^{k=1} - \xi_{\delta N=0}^{k=0}. \quad (6.3)$$

As a consequence, by subtracting every parabola by the first one ($k = 0$) and by rescaling with $\Delta\xi_0$, the manipulated ES displays some equally spaced structures with integer distances whose degeneracies are explained through a mapping with energy spectra of boundary conformal field theories [98]. These conformal structures are then destroyed after the phase transition to the gapped phase. We show the results for the XXZ and the BH models in Fig. 6.2 and Fig. 6.3.

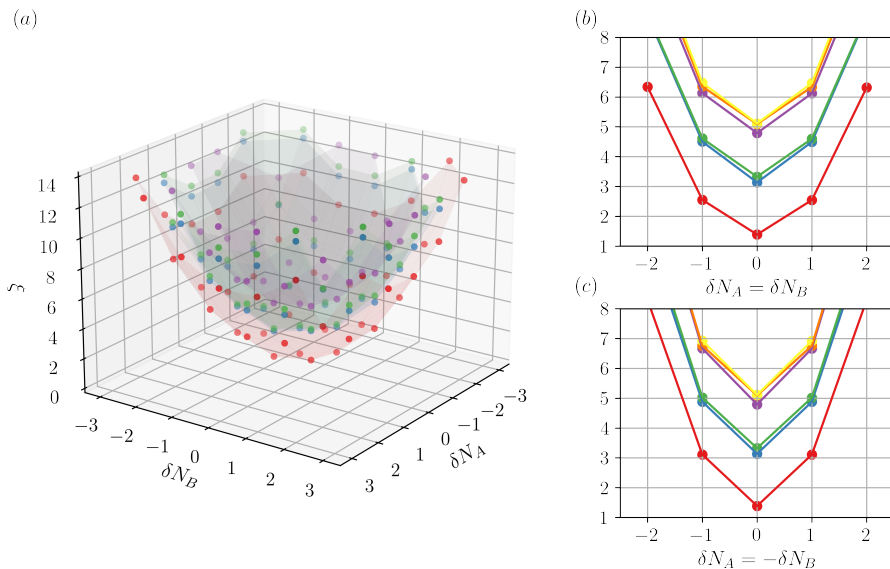


Figure 6.4: Entanglement spectrum of the $L = 64$ BH2S chain with periodic boundary conditions in the gapless phase for $U_{AB}/U = -0.15$. In (a) we show the parabolic surfaces on the $\delta N_A/\delta N_B$ plane only for the first four envelope for clarity. In (b) and (c) we plot the diagonal sections of the paraboloids corresponding to the density ($\delta N_A = \delta N_B$) and the spin channel ($\delta N_A = -\delta N_B$) respectively. After the phase transition to PSF, only the eigenvalues of the density channel (b) maintain the conformal behaviour while the spin channel acquires a gap as shown below in Fig. 6.5.

The presence of such parabolas (and their degeneracy counting) from a finite system is not guaranteed to be reliable though, since finite-size effect can spoil the conformal structures. For this reason, the ability of generalization typical of a machine driven approach is crucial for carrying out the pattern recognition for pursuing the distinction of the different regimes.




Figure 6.5: Entanglement spectrum of the $L = 64$ BH2S chain with periodic boundary conditions in the gapped PSF phase for $U_{AB}/U = -0.35$ in analogy to Fig. 6.4. In the spin channel (c) the envelopes are clearly no more parabolic.

More interesting properties are found in the generalization of the BH chain to the BH2S model where the situation is more complicated. The symmetries associated to the A and B number of particles' conservations provide two labels for the sectors δN_A and δN_B . The paraboloids become equally spaced paraboloids in the region 2SF phase (see Fig. 6.4). As already anticipated, once the system undergoes the phase transition to PSF the gap opens only in the spin channel. Indeed, as evident from Fig. 6.5 after the opening of the gap, the arrangement of the ES eigenvalues on the density channel slice of the paraboloids ($\delta N_A = \delta N_B$) still behaves according to the above-mentioned conformal structures while it loses these fascinating properties in the spin channel ($\delta N_A = -\delta N_B$). It is furthermore relevant that in the PSF the parabolic envelopes observed in the density channel are the only preserved, while the ES does not exhibit paraboloids in any other direction on the $\delta N_A/\delta N_B$ plane. We did not find any analysis about this generalization in the literature.

6.4 Further data analysis of the ES

In the previous section we have shown the peculiar structures appearing in the deep gapless and gapped phase, but we are not going to give as input to the




Figure 6.6: ES of the XXZ (left) and BH model (right) as a function of the control parameter. In the legend the first digit identifies the symmetry number (δN in the notation of the manuscript) and the second digit is the sorting index k .

GAN the symmetry resolved ES. In fact, the idea is to develop a generic procedure which should work regardless of the availability of the symmetric sectors. In such a case, the (low-lying) ES for a given configuration of the Hamiltonian parameters can be thought of as a vector in a high-dimensional space whose components are the eigenvalues. Since the transition is of the BKT universality class, we do not expect an abrupt change in the data patterns inbetween the deep phases, where the transition occurs. We show in Fig. 6.6 the evolution of several ξ_i for a 400 sites XXZ chain as a function of Δ and for a 256 sites BH chain as a function of U/t . We filter the sectors and plot only the ones with $\delta N \geq 0$ since the others have degenerate values due to particle-hole symmetry. In the figure, multiple crossings between the eigenvalues happen and one could in principle resort to the location of the crossings in order to perform a detection of the phase transition, in the spirit of level spectroscopy [106]. This is clearly a non-trivial task – considering the amount of points where crossings occur – which requires additional model-specific knowledge. Moreover, we show in the following that blindly considering the lowest-lying ES crossing (as it might seem somehow physical to do) can erroneously lead to the wrong extrapolation of the transition point. For the XXZ model it is worth mentioning that the additional $SU(2)$ symmetry at the transition point $\Delta = -1$ forces the ES to arrange into precise multiplets even for finite system sizes. This is not the case for the BH model in contrast. Actually, according to the deep analysis of Ref. [98], the transition to the MI is associated with a precise value of the ratio $\eta_C = \frac{\xi_1^0 - \xi_0^0}{\xi_0^1 - \xi_0^0} = \frac{1}{4}$. While this implies a series of crossings in the thermodynamic limit for the eigenvalues ξ_0^{k+1} and $\xi_{\pm 2}^k \forall k$ in correspondence to the critical value U_C , the same behaviour is not anymore guaranteed for finite sizes, as it can

be clearly seen in the right panel Fig. 6.6 for $k = 0$. Indeed, by looking at the left panel of Fig. 6.7 (i.e., the analogue of Fig. 3 of [98] with our data), the empirical η does not reach the critical η_C before the true transition point at finite size and the above mentioned crossings is avoided.




Figure 6.7: In the left panel the η ratio as a function of the interaction for the BH model for different system's sizes. In the right panel, three different extrapolations to the thermodynamic limit of the predicted critical coupling: the blue line according to the critical $\eta_C = 1/4$ method explained in the text, the brown line following the crossing between ξ_0^2 and $\xi_{\pm 2}^1$ and the green line following the first visible crossing between ξ_0^1 and $\xi_{\pm 1}^1$.

Even if the finite system does not contain the relevant crossing, in the left panel of Fig. 6.7 we show a possible extrapolation towards the thermodynamic limit of the pseudo-critical values $U_C(L)$ correspondent to $\eta_C = 1/4$ (in blue). Moreover we plot in brown the true relevant first available crossing between ξ_0^2 (the third eigenvalue in the $\delta N = 0$ sector) and $\xi_{\pm 2}^1$ (the second eigenvalue of the $\delta N = \pm 2$ sector). The two extrapolations nicely agree with each other and predict a critical value compatible with the one established in the literature (the shaded red area is computed considering the values from the literature as listed in [107]). Instead, following the lowest-lying ES crossing between ξ_0^1 and $\xi_{\pm 1}^1$ one would have obtained the wrong green scaling line in the right panel of Fig. 6.7.

We want to stress once again that our aim is to develop a method that is completely agnostic about the model-dependent value of η_C at the transition and about the symmetry number labelling of the eigenvalues which are both fundamental ingredients of the detection through the level spectroscopy.

From now on, we are going to treat the low-lying ES from a data science perspective: every configuration of the Hamiltonian parameter, is then associated

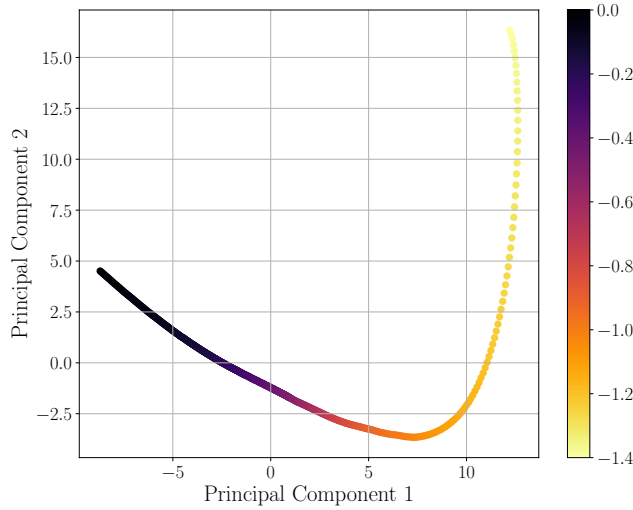


Figure 6.8: First two principal components of the PCA applied to the XXZ spectra for a 200 sites chain as a function of Δ (color dimension).

with a vector of eigenvalues. As we have seen in Fig. 6.6, its evolution as a function of the control parameter is rather complex and since the vector lives in a high-dimensional space we cannot easily visualize it. One possible solution is to resort to a dimensionality reduction algorithm such as the PCA [58] in order to come back to a low-dimensional representation. In Fig. 6.8 we show the first two principal components of the PCA – which explain approximately the 80% of the dataset variance – applied to the XXZ spectra for a 200 sites chain as a function of Δ . The selected dataset is composed of about 700 different spectra for values of Δ ranging between $-1.4J$ and 0. The color dimension represents the value of Δ . The two values of the two components already allow to separate between different patterns in the ES depending on the phase but, unfortunately, due to the nature of the BKT the data points do not cluster. The configurations in the two deep phases are, instead, continuously connected. In spite of that, the region of the transition appears to be located approximately at the elbow of the curve. Since the PCA is a linear transformation of the vectors, we expect that the employment of non linearities as in a NN can sensibly improve the ability of distinction.

6.5 Results with GAN

In order to carry out the BKT detection, we consider the biggest eigenvalues of the ES in magnitude for various system's sizes for all the three different models. We sort them in a decreasing way for the first ES corresponding to the origin of the phase diagram and we then maintain the same ordering so that the GAN is always fed with coherent levels. The GAN has no clue about the charge and ordering-index sector and therefore it has no a priori knowledge about the underlying parabolic structures. The maximum number of eigenvalues that was used, namely the input dimension for the GAN, is 64 since smaller eigenvalues are more easily affected by numerical error and, due to their magnitude, give a negligible contribution to the reconstruction loss.

We divide a portion of the data that we are sure belong to the gapless phase in two disjoint datasets: training set and validation set. Hence the training data show the above-described conformal structures. The validation points are instead taken from an interval of the parameter space adjacent to the training on the side toward the transition. With this splitting, we set a threshold both for the training reconstruction loss ($\bar{\mathcal{L}}_{rec}^{tr}$) and the validation loss ($\bar{\mathcal{L}}_{rec}^{val}$) at training time. The idea is that if we pick the validation points sufficiently far apart from the phase transition, we can check the ability of generalization of the network on unseen examples that we know belong to the same phase of the training set. In this way, we can be sure that the network is not overfitting the training data and that the reconstruction loss is not low due to the learning of statistical noise. Moreover, since in the first attempt we are training a different model for every system's size, we can make sure that the level of the losses are compatible and the comparison between the resulting loss curves is reasonable. With our GAN architecture we are able to achieve training losses of the 10^{-3} order within 250 epochs. It is fundamental to reach good values of precision for pursuing an accurate reconstruction of the ES's structures since the magnitude of the eigenvalues drops rapidly after the first leading ones. This level of precision is sufficient to faithfully reproduce at least the first two parabolas of the ES in the gapless region whereas the AEs become unable to reconstruct the structures after the phase transition. The training was performed by setting $\lambda = 0.1$ and $\epsilon = 10$ in Eq. (6.2). The typical threshold for the losses in the training and validation region for the XXZ model are $\bar{\mathcal{L}}_{rec}^{tr} = 0.00005$ and $\bar{\mathcal{L}}_{rec}^{val} = 0.0001$, for the BH model are $\bar{\mathcal{L}}_{rec}^{tr} = 0.005$ and $\bar{\mathcal{L}}_{rec}^{val} = 0.02$ and for the BH2S model are $\bar{\mathcal{L}}_{rec}^{tr} = 0.005$ and $\bar{\mathcal{L}}_{rec}^{val} = 0.05$. The parameters of the network




Figure 6.9: The reconstruction loss percentages as a function of Δ/J for different sizes of the systems (indicated in the legends). The location of the BKT transition in the thermodynamic limit is reported with the dashed red line. The training set include points in the yellow interval whereas the validation samples are taken from the green interval. All the curves show an abrupt rise in the proximity of the transition.

were optimized through ADAM optimizer with learning rates of the generator $\text{lr}_G = 0.01$ and of the discriminator $\text{lr}_D = 0.0001$. For both optimizers, a `CosineAnnealingLR` scheduler (see in documentation of [94]) was employed in order to adjust the learning rate at training-time and improve the convergence. In this first part we provide the results for all the three models and we discuss the differences between them: once again, for every size of a given system a different GAN architecture was trained and evaluated. The best scenario would be to obtain a zero loss (perfect reconstruction of the GAN) in the training phase and very high loss in all the other phases, hence emulating a fictitious order parameter of the transition. However, the transition for the finite system is not sharp, especially for the BKT universality class. We expect a much more evident jump in the loss for other transitions, like Ising or first order.

We show the results for the XXZ model in Fig. 6.9 for which we have trained the GANs with samples taken from the region $-0.65 \leq \Delta/J \leq 0$ (about 550 points included in the yellow shaded region of the plot) and validated with configurations in $-0.8 \leq \Delta/J < -0.65$ (green shaded area). The loss rises near the phase transition at $\Delta = -1J$, indicated with the red dashed line. The enhancement of the change in the loss respects the expected behaviour as a function of the system's size, toward the thermodynamic limit. As a difference with respect to the previous works involving the same scheme [73,




Figure 6.10: The reconstruction loss percentages as a function of U/t for different sizes of the systems (indicated in the legends). The location of the BKT transition in the thermodynamic limit as numerically estimated in the studies listed in [107] is represented as a red shaded region. In this case the loss starts to rise already in the validation region but remains under the 10%, instead in the region of the transition the behaviour of the curves changes according to the size of the system. Even by changing the training window (see below), the splitting of the curves remains very closed to the transition.

[108] in which the authors exploit tensor network techniques for the simulation of the thermodynamic limit, we cannot expect a sudden change in the ES structures along the transition.

The BH training consists in about 500 points from the interval $0 \leq U/t \leq 2.5$ and the validation from $2.5 < U/t \leq 3$. The results for the model are represented in Fig. 6.10. In this case the location of the transition can only be computed numerically. We show in the figure the estimate of the critical value as a shaded region that is computed considering the values from the literature (as listed in the Table 1 of [107]) and providing an interval centered in their mean and large as the standard deviation. The losses of the GANs start to increase already just outside the training set (remaining below 10% though) but the sharp rise occurs near the latter region. Quite peculiarly, the curves of different system's sizes begin to separate from each other in the vicinity of the critical point.

Finally, for the BH2S model there are about 200 samples from the region $-0.1 \leq U_{AB}/U \leq 0$ as training points and the validation ones are from $-0.15 \leq U_{AB}/U < -0.1$. The results for the model are represented in Fig. 6.11. For this model, there is no analytical prediction of the critical value and the location




Figure 6.11: The reconstruction loss percentages as a function of U_{AB}/U for different sizes of the systems (indicated in the legends). The predicted value for the BKT is still uncertain. The plot shows a trend very similar to the single species BH model. We hypothesize that the transition should happen in the $-0.25 < U_{AB}/U < -0.2$ interval where the curves begin to separate in analogy to the previous case. This result is in good agreement with the results of the study of the $R_S(x)$ correlation.

of the transition is still uncertain. In the Supplemental Material of [63] (see also references therein) we detect the PSF transition from the curvature of the fitting function of the $R_S(x)$ correlation in the log-log plane (as in Fig. 5.4 panel (c)): since the latter correlation decays exponentially in the PSF phase and, instead, algebraically in the 2SF phase, the curvature is expected to be zero in the latter case. Therefore the fit is a line in the log-log plane representation of R_S while the curvature starts to be finite in the PSF regime and the line bends.




Figure 6.12: The values of α for different system's sizes: $L = 8, 16, 32, 64, 96$ (from light to dark colors). An increase in the α amplitude is associated to an exponential trend of the R_S correlation function.




Figure 6.13: The loss curves for the XXZ with $L = 400$ sites in (a) and for the BH model with $L = 256$ sites in (b) for different training regions indicated in the legend and accordingly to the curves' color by the shaded regions. The KL divergence is represented with the black dashed line, whose values must be consulted in the rightmost y-axis.

We report the calculations of the curvature, named α also here in Fig. 6.12. Despite the small resolution in the U_{AB} values, the curves cross between each other for values of $-0.25 \leq U_{AB}/U < -0.2$. Since the losses in Fig. 6.11 show a trend very similar to the single species BH model, if we hypothesize that the transition should happen where the curves of different sizes begin to separate, the agreement with the curvature method will be very good.

In order to check the stability of the algorithm, we verified that the results do not show substantial differences when the training intervals are rescaled. We report in Fig. 6.13 the results obtained by varying the width of the training window in the parameter space for the XXZ and the BH model. The loss computed from the GAN reconstruction is also compared with the Kullback-Leibler divergence (KL) defined as follows:

$$D_{KL}(P||Q) = \sum_i P(x_i) \log \left(\frac{P(x_i)}{Q(x_i)} \right), \quad (6.4)$$

which is a measure of the relative entropy between two probability distributions and therefore is a proper measure for the distance between two ES. The reference distribution was taken to be the ES at the origin of the phase diagram and the distance from it was computed for every ES labelled with the control variable. It is represented in Fig. 6.13 as a black dashed line whose values run over the rightmost y-axis. For the XXZ case the KL has a delayed and slow change after the transition while for the BH model the change occurs in the right region but with a much less sharp rise than the loss performed with the

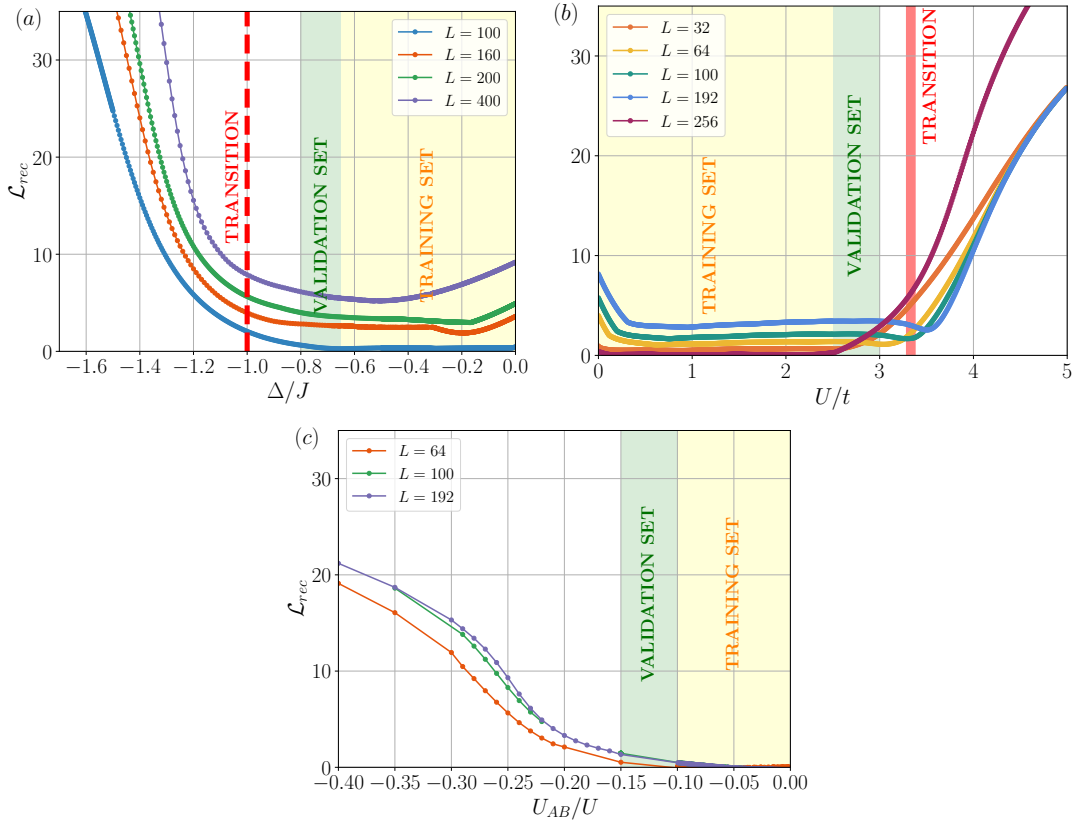


Figure 6.14: The reconstruction loss percentages for the three models when the GAN is trained the single smallest system’s size and then evaluated on the datasets of all the other sizes. The training and validation regions are the same as in Figs. 6.9, 6.10 and 6.11. The BH2S is actually trained on the 64 sites system because we observe a quite mediocre performance in the reconstruction for the model trained on that size, probably due to the fact that the system is too small and strong finite size effects play a crucial role.

GAN. On the one hand, the KL divergence line shows how the measure of the distance between representative quantities of the ground state in two different phases provides an interesting measure to observe changes due to the phase transition, in the spirit of ground state fidelity [109]. On the other hand, it is not sufficient. The GAN flexibility through the anomaly detection scheme allows, instead, to have at disposal a more powerful tool that can automatically find the common factor among different patterns within one phase and differentiate them from the patterns of the other regimes.

We obtain very similar results to Figures 6.9, 6.10 and 6.11 when training only one GAN per model on a single system’s size and then do the evaluation on the datasets of all the other sizes. This procedure could be extremely useful when the computational cost of producing a training set is particularly high and a




Figure 6.15: Loss curves for the BH model when the GAN is trained in the MI phase. The training and validation regions are the yellow and green windows respectively. Even though the GANs are probably overfitting the very limited datasets, the behaviour of the loss in the anomalous phase is quite interesting.

rough estimation of the phase diagram boundaries is enough in a first approach. We represent in Fig. 6.14 the loss curves for the XXZ, the BH and the BH2S models when the GAN is trained on a single system's size and then evaluated on the datasets of all the other sizes. The training and validation regions are the same as in the previous figures. Even though the loss is not really low in the training phase – anyway remaining below 10% –, the increase happens again near the transition. Quite interestingly, the loss also goes up approaching the origin of the phase diagram ($\Delta = 0$ and $U = 0$ respectively): presumably it is due to the fact that those are singular points where the absence of anisotropy and repulsive interaction makes the ground state possess different properties.

Since the data points in the gapped phase are not as numerous as for the gapless one, we cannot perform an equally detailed analysis with the anomaly detection protocol with a training on that side. However, we can still check the behaviour of the loss function in the region of the transition. We report in Fig. 6.15 the results obtained for the BH model when the GAN is trained in the MI phase. The typical threshold for the losses in the training and validation region were higher than before, with $\bar{\mathcal{L}}_{rec}^{tr} = 0.008$ and $\bar{\mathcal{L}}_{rec}^{val} = 0.07$ respectively. Despite the rise of the curves seems to start earlier than the expected transition, it is worth mentioning that their behaviour in the gapless region is quite interesting: after the sharp increase the loss stabilizes meaning that, even if the GAN is not able to reconstruct the ES in the gapless phase, the level of error remains almost constant for abnormal examples within another phase. We attribute




Figure 6.16: Loss curve for the BH model ($L = 256$) when the GAN is trained in the transition region (yellow shaded area). The loss increase outside the training region remains below 10%, meaning that the network has learnt features of both the gapless and gapped phase.

the worse performance to the very limited number of training examples (about 100 samples per system's size) that probably results in overfitting networks (even making use of regularizations) and poor generalization.

Lastly, another natural question about the anomaly detection scheme is what happens if one does not know the arrangement of the phases at all and tries to train the model with configuration points coming from the transition region. In this case we observe a completely different loss profile as in Fig. 6.16 for the $L = 256$ BH chain. The loss increase outside the training region is very restrained and always stays below 10%. We argue that the GAN, being trained on data coming from both the phases, has learnt how to reconstruct quite faithfully the ES in both the cases and therefore it does not identify any anomalies. Even without any knowledge of the phase diagram, it is possible to find the proper regions where to train by iteratively scanning several training windows and looking at the resulting loss profiles.

6.6 Loss threshold and precise transition detection

At this point the missing piece for completing the anomaly detection protocol is the definition of a threshold for the loss function in order to state what are

the configurations belonging to the same phase of the training points and what are the ones of other phases. Setting a threshold by hand is not a systematic criterion, hence we propose to look at the knee of the loss curves, defined as the point where the loss curve has the maximum curvature, meaning that the change rate is the maximum possible. In the thermodynamic limit, we expect this point to approach the true critical value since the loss increase should become sharp. We report in Fig. 6.17 the knees of the loss curves for the XXZ, the BH and the BH2S models. The knee position was computed by resorting to the library described in [110], where the authors provide a general approach to on line and off line knee detection that is applicable to a wide range of systems. The knee position is formally defined for continuous functions and the comparison of the results of the method on discrete sets of data is very well explained and motivated. The location of the knees are reported as vertical dashed lines using the same color code as the dimension of the system in Fig. 6.17: the XXZ shows a quite slow convergence of the knee position toward the transition while the BH has a very good agreement with the literature results already for small system sizes. For the BH2S model we only report the knee position of the biggest system ($L = 128$) because we miss some points for the other sizes. The predicted pseudo-critical point is quite near the region $-0.25 < U_{AB}/U < -0.2$ estimated with the α coefficients of the correlation function fit but clearly more resolution would be needed for a more consistent comparison.

Despite the results look very promising, primarily because the critical positions are quite near the expected values even for finite (and sometimes small) systems, we think it is worth to keep in mind the main purpose of the anomaly detection scheme. We believe that the machine learning approach to the phase diagram mapping should not be considered as an alternative to the physically well-motivated traditional techniques studying the order parameter (when present), the gap and so on. The real use case of these method should be as a complementary tool to the traditional techniques when the phase arrangement is not clear or when the order parameter is not well defined, or again when an automatic procedure is needed to roughly (and if used carefully even quite precisely) draw the phase boundaries.

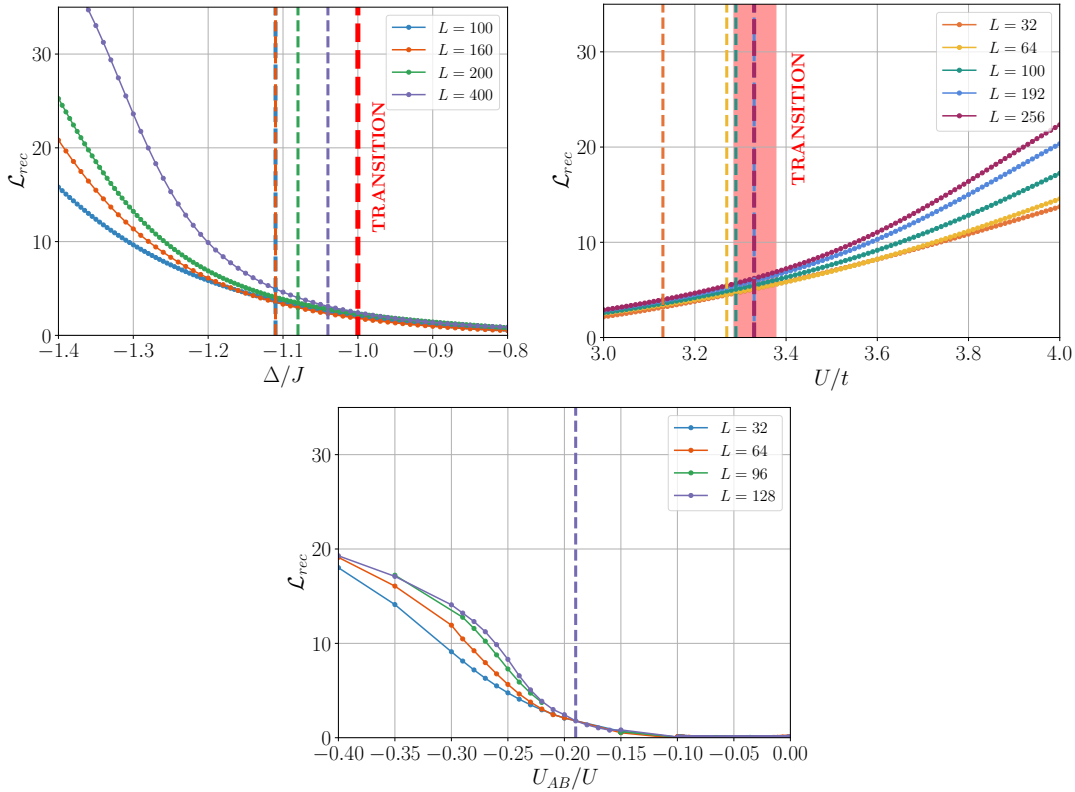


Figure 6.17: Zoom on the transition regions for the same loss curves as in Fig. 6.9, 6.10 and 6.11. The vertical dashed lines represent the knee of the curves position as computed with the method described in [110]. For the anomaly detection protocol applied on finite systems, we believe that a loss threshold systematically set on the curve knee can provide a good estimate of the pseudo critical point. Notice that for the BH2S model only the knee of the $L = 128$ curve is reported because all the other would lie deeper in the gapless side (in analogy to the BH) but we lack some points to obtain an acceptable resolution in the coupling U_{AB} .

6.7 Conclusions

In this section we show that the ES represents a reliable quantity to look at in order to perform a machine-driven detection of the elusive BKT transition for one-dimensional quantum systems by assuming almost no prior knowledge about them. We face such a task by using a GAN that benefits from the adversarial training in an anomaly detection scheme. We find that such an architecture is a valid candidate for the purpose of discerning the change of the patterns of the ES. We support our claim by presenting the results of the application of the anomaly detection scheme on numerical data for the XXZ, BH and BH2S models. Although we discuss the results obtained with

the analysis of the ES, our method is very general and can be applied to any representative quantity of the many-body state that embeds the changes in the physics associated to the phase transition. For all the three models we show that the results are quite promising and the GAN can give a good estimate of the critical point which in some cases agrees with the literature results even for quite small system sizes. Let us conclude by remarking once again, that this approach is not supposed to become a substitute to the traditional methods for the phase transitions detection but allows to obtain a qualitative map of a phase diagram with very little knowledge about the nature and arrangement of the phases in an automated fashion.

Chapter 7

Gaussian fit of PDF

In the previous sections 6.3 and 6.4 we have shown that the properties of the bipartite ES represent a powerful indicator of the many-body wave-function properties, specifically for the purpose of characterizing the quantum phase transitions. Actually, in recent years, several other quantities of bipartite systems have emerged within the quantum information and condensed matter theory cross-fertilization. With their help, quantum phases of matter between them have been studied, and put in relation with the underlying quantum field theories. A first prominent example is the log-scaling of the von Neumann (and, more generally, of any R enyi) entanglement entropy with the bipartition size for one-dimensional critical systems with local Hamiltonians [12, 16, 111]: fitting the coefficient in front of such law is arguably one of the best ways to estimate the so-called central charge c of the associated conformal field theory (CFT). As a counterpart, we have already described the features of the full ES both in gapless and gapped phases but several other characteristics have been studied in the literature starting from the degeneracy patterns of topological phases [23, 100], the embedded information about non-local quantum correlations [101, 102, 103] and the relation between specific eigenvalues for the definition of ad-hoc order parameters [104, 105]. Indubitably its wealth of structures containing properties of the physics of the system makes the ES a very suiting tool for machine learning approaches [68, 73] for the phase transitions' detection task. Last, but certainly not least, the evolution of entanglement quantities under the system dynamics has recently unveiled the existence of new kinds of non-equilibrium phase transitions for quantum many-body systems under random projective measurements or unitary gates [112, 113].

As a consequence of the above, considerable efforts have been spent towards making such entanglement features measurable in the laboratory [114, 115, 116, 117, 118]. Among the many, the so-called number entanglement entropy is gaining a prominent role: its operational definition refers to the probability density function (PDF) of a U(1) globally conserved charge in an extensive sub-portion of the system, also for mixed states [119]. Evidences of its distinctive dynamical behaviour as a hallmark for many-body localization have been recently obtained both numerically and experimentally [120, 121, 122].

Some specific properties of the PDF have been explored in the past, above all its second momentum or, in other words, the amount of charge fluctuations \mathcal{F} across sub-systems. As explained in [123], a lot of properties of the fluctuations are shared with the entanglement entropy: for a gapped phase, they exhibit a strict area-law behaviour, $\mathcal{F} \propto L^{d-1}$, with L the linear size of the partition and d the spatial dimension of the system, whereas for a gapless phase there appears a logarithmic correction: $\mathcal{F} \propto L^{d-1} \ln L$ [124]. In particular for a Luttinger liquid, the scaling coefficient is related to the K -parameter, thus yielding yet another piece of information about the underlying field theory (in the Appendix A we present the results of such a scaling analysis for the fitting of K in a specific case).

We propose in the following a new approach to map out the phase diagram of quantum many-body systems by considering the full PDF. We are able to detect all the phase transitions of the one-dimensional extended Bose-Hubbard (EBH) model at zero temperature by performing a simple and yet very general procedure consisting of some educated fits of the PDF only, therefore without resorting to any phase-specific quantity like order parameter or correlation. We show how the PDF, being intimately related with the ES, preserves its intrinsic abundance of information about the nature of the phases. This can be exploited for an agnostic and automatic detection, as done with the machine learning solutions to the problem like the one we presented before. The huge advantage of the full PDF is its availability in experiments, e.g. with quantum gas microscopes, and not only in numerical simulations like the ones we perform here via MPS. We also show that a finite-size scaling analysis of the PDF leads to a pretty precise determination of the phase boundaries and that some previously found functional forms of the PDF for limiting scenarios [111, 125, 126, 127, 128, 129] can be connected to each other.

7.1 Number probability density function in the EBH model

We consider the ground state of the EBH model as divided in two equal complementary portions A and B . While the total number of particles is conserved, the number in the single region can fluctuate. The measurement of a deviation $\delta N = m$ from the average density happens with a probability:

$$p(\delta N = m) = \text{Tr}(\rho_A^{\text{red}} \Pi_{N/2+m}) = \sum_i p_i^{(m)}, \quad (7.1)$$

where $\Pi_{N/2+m}$ is the projector on the sector with $N/2 + m$ occupancy. We introduced the eigenvalues p_i of the reduced density matrix ρ_A^{red} with an additional label m referring to the imbalance in the number of particles in the A region: $N_A = N/2 + m$. Once again, being the $p_i^{(m)}$ s the natural metric on which truncations of the Tensor Network representation are performed, the PDF is automatically at disposal without extra computational costs. It is worth noticing that in an experimental setup with access to site-resolved populations, the PDF is obtained by bin-counting the occupation numbers in half of the system [120].

For gapless one-dimensional phases described by a CFT, the PDF is a Gaussian,

$$p(m) \propto \exp(-\beta m^2), \quad (7.2)$$

due to the presence of a free bosonic generator in the theory [123, 126, 130]. This is also associated to the equally spaced parabolas of the ES which were described before: $\xi_m^k = \xi_0^0 + k\Delta\xi_0 + \beta m^2$, where the sorting index k denotes also the order of the parabola (some could be degenerate), ξ_0^0 is the lowest eigenvalue of the ES in the $m = 0$ sector, $\Delta\xi_0$ the difference between the lowest and the second eigenvalue in the $m = 0$ particle sector [98] as defined in Eq. 6.3. As we will detail below, we expect instead the PDF of gapped phases to exhibit sensible deviations from this picture [99].

From the operational perspective, the differences between Gaussian and other PDF profiles can be captured via the average residuals of a two-parameters linear fit, namely

$$\text{Res} = 1/N_m \sum_m (\log p(m) + \beta m^2 + \beta')^2, \quad (7.3)$$




Figure 7.1: Phase diagram of the EBH model traced with the residuals of the PDF's quadratic fit $p(m) \propto e^{-\beta m^2}$ for a $L = 64$ system. Details about the phase transitions along the dashed and dotted cuts are shown in Fig. 7.4. The PDFs correspondent to the configurations marked with the colored shapes are displayed below in Fig. 7.5.

with N_m the number of fitted points and β' a normalization-related parameter. In Fig. 7.1 we plot the logarithm of such quantity over the span $m \in [-4, 4]$ for a fit performed upon the distribution values for $m \in [-2, 2]$.

This simple procedure turns out to be very effective to map out the entire phase diagram even without prior knowledge of the phases and the PDF shapes to be expected. We find perfect agreement with previous studies [73, 84], without resorting to any investigation of the correlation functions or gap nor to the use of machine learning techniques and – even more importantly – by using an experimentally accessible quantity.

In this respect, in analogy with Fig. 7.1, we reproduce the phase diagram of the EBH model by looking at the residuals of the Gaussian fit on data of a simulated experiment. Instead of considering the numerical PDF of each pair of couplings (U, V) , we sample a given number of shots N_S from it. Every shot corresponds to a value for the imbalance in the number of particles m among the two subsystems. Once the samples are extracted, we fit the histogram of the sampled distribution and compute the residuals from there, therefore emulating an actual experiment that deals with the counting of the particles from the snapshots of the system. The results are very promising, as shown in Fig. 7.2. For a low number of samples ($N_S = 100$ and $N_S = 500$, top row) the only distinction between the deep MI and CDW phases is possible. The phase separated SS-SF is already well distinguishable for the reasons explained




Figure 7.2: Phase diagram as of Fig. 7.1 obtained from different simulated experiments: the PDF is fitted from 100,500,1000,10000 shots (from left to right, top to bottom).

in the following. Once one proceeds up to $N_S = 10,000$ shots (bottom right), the level of detail is good enough to even appreciate the HI and, in general, to be comparable to a numerical study. Notice that those numbers of snapshots are normally achievable in state-of-the-art experiments with ultracold atoms on lattice since the current repetition time of an experiment is of the order of seconds. An hypothetical strategy could be to roughly draw the phase diagram with a few samples for every (U, V) on a coarse grid and then refine the phases' borders with a greater (nonetheless feasible) number.

Both in Fig. 7.1 and in Fig. 7.2 even with only 100 samples, our method properly identifies the recently postulated SF+SS phase, which appears as a very noisy region for large V and small U . This is a consequence of the alternation of uniform density regions and density-wave ordered ones: the PDF is peaked on a value which is not necessarily the average density but depends on the location of the supersolid domains. Since the PBC let these domains emerge in random positions along the ring, this gives rise to the above-mentioned noise. In order to better grasp the situation in this phase, we plot in Fig. 7.3 the PDF for $(U, V) = (0.5, 4)$ of a system of length $L = 64$ sites. The PDF is clearly not centered in $m = 0$ in contrast with all the other discussed phases. Indeed, by looking at the density profile of the system in the inset, a supersolid domain




Figure 7.3: In the main plot the PDF of a configuration in the phase separated regime SF + SS for a system of $L = 64$. The inset shows the correspondent profile of the density.

has emerged with a periodic structure in the left region of the lattice while the right region is flat and with a superfluid character. A very detailed explanation about the phase separation can be found in [131].

The nature of the phase transitions is reflected in the evolution of the residuals in their vicinity: in Fig. 7.4 we show the trend of the average residuals for two cuts (also indicated in Fig.7.1 with dotted and dashed lines) of the phase diagram and different system's sizes ($L = 16, 32, 64, 128, 256$ from light to dark color). The first cut of Fig. 7.4 (a) contains two gapped-to-gapped phase transitions for fixed $U = 5$, from MI to HI and from HI to CDW as indicated with a dotted line in Fig. 7.1. The location of the critical points is easily determined by the abrupt decrease of the residuals which becomes sharper and sharper as the length of the ring increases: in the two insets Fig. 7.4 (b-c) a detail on the scaling of the critical V_C against the inverse of the system size is shown together with a linear fit. The extrapolation of the critical values is in perfect agreement with previous studies, i.e., $V = 2.95 \pm 0.05$ for the transition MI-HI and $V = 3.525 \pm 0.05$ for HI-CDW [84].

The second cut Fig. 7.4 (d) encompasses the gapless-to-gapped BKT phase transition from the SF to the MI at $V = 0$. The plot is very similar to the one obtained with the GAN in the previous chapter. It would be interesting to find a scaling procedure for the residuals, similarly to what is performed for the superfluid stiffness and/or the K -Luttinger parameter in standard approaches: At the moment, this remains an open problem. We stress, however, that the

β parameter of the linear fit offers a direct experimental access to the K -Luttinger parameter [123]. For completeness, we describe the results of such an analysis in Appendix A (see, also, Ref.[77, 78] therein) and we obtain a critical value $U_c \simeq 3.36 \pm 0.01$ in perfect agreement with previous studies as reported in [107]. From the Gaussian fit perspective, the location of this transition corresponds to the uprising of the residuals.

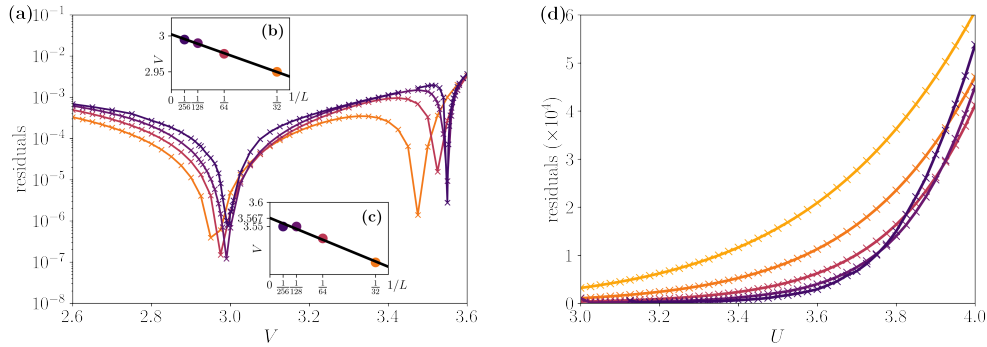


Figure 7.4: Residuals of a Gaussian fit of the PDF $p(m) \propto e^{-\beta m^2}$ for different system's sizes $L = 16, 32, 64, 128, 256$ from light to dark color in the proximity of (a) the gapped-to-gapped phase transitions MI-HI-CDW along the cut $U = 5$ and (d) the gapless-to-gapped phase transitions SF-MI (BKT) for $V = 0$. In the insets, the location of the residuals' minimum is fitted against the inverse of the system's size in order to extrapolate the critical V_C for the MI-HI phase transition (b) and the HI-CDW transition (c).

7.2 PDF for gapped phases

So far we presented the results of a pure Gaussian fit of the PDF as a tool for the mapping of the phase diagram. The procedure relies on the presence of a gapless phase where the PDF is Gaussian. Actually, a more detailed look at the PDF unveils more information than the simple residuals to Gaussian fitting discussed above. In Fig. 7.5 we present some prototypical configurations in the different phases, corresponding to the coloured symbols in the phase diagram of Fig. 7.1: e.g., panel (a) shows the perfect parabola for the SF phase ($(U, V) = (0.5, 0.5)$). The striking feature is that all PDF profiles can be captured by proper shifts and combination of Gaussian envelopes, as we explain here below.

The shape of the PDF deep in the gapped phases can be computed resorting to boundary-linked perturbation theory [99]. For example, the ES lev-




Figure 7.5: (color online). The logarithm of the PDF (black points) is plotted for a representative configuration of the **(a)** SF phase (red star, $(U, V) = (0.5, 0.5)$), **(b)** MI phase (green pentagon, $(U, V) = (5.5, 0.5)$), **(c)** CDW phase (gray square, $(U, V) = (2, 4.5)$), and **(d)** HI phase (brown diamond, $(U, V) = (5, 3.3)$). The gray lines show the fit of the PDF in terms of Gaussian envelopes and their shifts, as discussed in the main text. For both MI and CDW, the behaviour is linear for small $|m|$, in compliance with horizontally shifted parabolas: on the right we show the shift of parabolas' minima from $m = 0$ as a function of the inverse distance from the critical point for MI **(e)** and CDW **(f)**. It is apparent that they tend respectively to 1 and 2 deep in the phases, as predicted by perturbation theory in the text. For the HI, the dominant feature is a vertical shift between parabolas for the different parities of m : **(g)** shows a magnification thereof.

els for the MI are obtained by consecutive applications of the kinetic term $H' = \sum_{j=1}^L -t(b_{j+1}^\dagger b_j + \text{h.c.})$ on the zero-order ground state, i.e., $|\dots 111\dots\rangle$. We are then interested primarily in achieving a given inbalance m with the minimal number of moves (i.e., perturbative orders of H'). In the case of a single boundary, it is rather easy to see that the leading order amounts to $(t/U)^{|m|(m+1)/2}$, up to a global weight depending on the ratio between the accumulated bosonic factors and the excitation energies along all possible sequences of moves.

In particular, by following the standard prescription with a perturbation $H' = \lambda W$ to the original hamiltonian, the relevant states at k -th order in our case are:

$$|\beta^{(k)}\rangle \propto \lambda^k \frac{W_{\beta l_1} W_{l_1 l_2} \dots W_{l_{k-3} l_{k-2}} W_{l_{k-2} \alpha}}{E_{\beta \alpha} E_{l_1 \alpha} \dots E_{l_{k-2} \alpha}} |\alpha^{(0)}\rangle \quad (7.4)$$

using the notation for the perturbation matrix elements and energies:

$$\begin{aligned} W_{\beta\alpha} &\equiv \langle \beta^{(0)} | W | \alpha^{(0)} \rangle, \\ E_{\beta\alpha} &\equiv E_{\beta}^{(0)} - E_{\alpha}^{(0)}. \end{aligned} \tag{7.5}$$

All the other terms are sub-leading because, for fixed m , require more than the minimum number of moves in order to transport the particles from subsystem B to A (and vice-versa for negative m) and hence belong to a higher perturbative order. The recipe for the calculation of the full corrections, involves the summation of all possible combinations for moving m particles though. Such a summation must be weighted with the ratio between the accumulated bosonic factors resulting from $\langle \beta | \sum_{j=1}^L -t(b_{j+1}^\dagger b_j + \text{h.c.}) | \alpha \rangle$ and the energy. For the single boundary case, by following the above mentioned protocol one finds that the amplitude of the states is proportional to $(t/U)^{\frac{|m|}{2}(|m|+1)}$ for the MI starting from the ground state $|1, 1, \dots, 1, 1\rangle$ (in agreement with [99]). When dealing with two boundaries, simple combinatorics leads to $p(m) \simeq \log(t/U) \left\lfloor \frac{(|m|+1)^2}{4} \right\rfloor$. The expression describes a symmetric envelope around the average number of particles, which could be recovered also by considering two symmetrical Gaussians shifted by ∓ 1 with respect to $m = 0$, for the positive and negative values of m , respectively. Remarkably, we find such horizontally shifted envelopes to persist with reduced offset when approaching the critical point, finally merging back when transitioning to the superfluid phase, see Fig. 7.5 (b-e).

We also highlight here that a similar shift is the dominating feature of the PDF for the CDW phase (dark gray square for $(U, V) = (2, 4.5)$), this time tending to ± 2 deep in the perturbative regime, hinting at the underlying structure of the zero-order ground state, i.e., $|\dots 0202\dots\rangle$, see panel 7.5(c-f).

The appearance of the $|m|$ dependence in the exponent due to a non-zero shift of these double parabolic envelopes can be exploited to distinguish such gapped phases from all the others. We verified that the values of the α coefficient of the fit:

$$-\log[p(m)] = \alpha|m| + \beta m^2 + \gamma \tag{7.6}$$

is non-zero only in the MI and CDW. The change rate of the coefficient along the phase transitions manifests again their nature as represented in Fig. 7.6. In the latter, we plot the value of the α coefficient as a function of the couplings of the model, in analogy to Fig. 7.1. The non-zero value of this coefficient




Figure 7.6: Linear dependence of the logarithm of the PDF for small m as a function of the interactions.

is the reason why a purely quadratic centered fit $p(m) \propto e^{-\beta m^2}$ fails in the above mentioned phases, giving rise to big residuals. For the single boundary partition of some specific models, instead, the PDF may assume asymmetric shapes, see a very nice description for the XXZ open chain in [99].

Finally, a direct inspection of the PDF in the HI phase (brown diamond $(U, V) = (5, 3.3)$) reveals the typical footprint of the topological order. A detailed description can be obtained by truncating the maximum site occupation to $n_{max} = 2$ bosons and mapping the EBH model to a spin-1 Heisenberg model [132, 133, 134]. In this framework the configurations of the HI appear as a *dilute anti-ferromagnet* of doblons(2) and holons(0) separated by an undetermined number of single occupations(1), $|\dots 21 \dots 10 \dots 21 \dots 10 \dots\rangle$ [133]. Noticeably, the PDF of this topologically gapped phase, panel 7.5 (d), resembles very closely the Gaussian PDF of the gapless phase, with the only deviation of a slight shift for even/odd Gaussians (see Fig. 7.5 (g)), alluding to the underlying parity-string order parameter. This explains the sensibly reduced, though still sizeable, residuals in Fig. 7.1.

The inclusion of the linear term via the α coefficient in the fit can give additional insight in the phase diagram mapping. By performing a fit of the form of Eq. (7.6) to the PDF and computing the residuals of the fit, we can faithfully reproduce the PDF of the SF, MI and CDW phases, thus leaving only the SS+SF and HI regions. We plot the results of the residuals in Fig. 7.7.




Figure 7.7: Phase diagram of the EBH model drawn with the residuals of the PDF's fit including the linear term for a $L = 64$ system. The SS+SF and HI regions are identified by the high value of the residuals.

7.3 Conclusions

In conclusion, we have shown that the probability density function of the occupation number of a portion of the system (and simple fits thereof) can be a powerful agnostic inspection tool for the phase diagram of quantum many-body systems. For the extended Bose-Hubbard model, the obtained results are comparable with much more sophisticated analysis both via traditional methods dealing with gap scalings and correlation functions [84] and via modern machine learning approaches fed with the ES [73]. We claim that the PDF provides the best intermediate quantity between the whole ES and the bipartite number fluctuations [111] also for automatic detection protocols, since it combines the advantage of being easily accessible in modern experiments (e.g., quantum gas microscopes) with the wealth of information about the full many-body state, while requiring little prior knowledge about the emerging phases. Moreover, we foresee the method to be valid more in general: an extension beyond zero-temperature regimes and one-dimensional systems will be an appealing follow-up of this work.

Chapter 8

Feature selection

This chapter is devoted to the applications of a particular kind of autoencoder which can be used to perform an automatic *feature selection* protocol. The theme of feature selection is quite important and very timely in deep learning for several reasons: in our case the goal is to select automatically the most relevant features for the problem of detecting phase transitions and thus remove the irrelevant or redundant ones so that the model can predict with increased accuracy; moreover, feature selection allows to make the model interpretable, therefore helping to understand the underlying patterns in the data. The architecture that we use is called Concrete Autoencoder (CAE) and it is inspired by the work [135]. We will present here the results of the study made with the CAE for the detection of the BKT in the BH model using the entanglement spectrum as input data. We aim at selecting the most important eigenvalues of the ES to be used for such a detection task.

In section 3.2 we have described how the AE architecture is implemented to learn a compression function $f_\phi : \mathbb{R}^d \rightarrow \mathbb{R}^k$ with $k < d$ and a decompression $g_\theta : \mathbb{R}^k \rightarrow \mathbb{R}^d$ with known shape but unknown parameters θ and ϕ , s.t. the reconstruction of the input \vec{x}' is a copy of \vec{x} . During the training, the AE's weights are optimized so that they minimize the empirical reconstruction error $d(\vec{x}, \vec{x}')$ over the training dataset $\mathcal{D} = \{\vec{x}_1, \dots, \vec{x}_N\}$ of i.i.d samples of some unknown $p(\vec{x})$.

Now we use a very peculiar compression function f_ϕ that selects a subset of the input features \vec{x} belonging to \mathcal{S} s.t. $\dim(\mathcal{S}) = k < d$ from which the input can be entirely reconstructed. This task is appropriate for those datasets where

lots of features are correlated to a minimal subset of independent ones. We do not want to specify some a priori choice for the subset though, instead, we want the network to learn automatically what is the best subset of features so that the reconstruction error is minimized. The only free parameter is then the hidden dimension k in the bottleneck of the AE. This filtering procedure is a kind of automatic feature selection and can be pursued with a special trainable mask to be applied to the input. The latter can be formalized as a matrix multiplication:

$$f_\phi(\vec{x}) = M_\phi \cdot \vec{x}, \quad (8.1)$$

where M_ϕ has d columns and “one-hot” rows labeled with i , $\vec{m}_j^{(i)} = 0, 1$ and $\sum_{j=1}^d m_j^{(i)} = 1 \forall i$ so that $f_\phi(\vec{x}) = \vec{u} \in \mathbb{R}^k$ whose components are a subset of \vec{x} components. The M_ϕ dependence on the ϕ s must be differentiable so they can be eligible to learnable parameters. The solution discussed in [135] consists in the sampling of the mask from a *concrete distribution* over the parameters $\phi \in (0, \infty)$ that is implemented as follows:

$$m_j^{(i)}(\phi, T) \sim \frac{\exp[(\log \phi_j^{(i)} - g_j^{(i)})/T]}{\sum_{l=1}^d \exp[(\log \phi_l^{(i)} - g_l^{(i)})/T]} \equiv \text{Softmax}[(\log \phi_j^{(i)} - g_j^{(i)})/T], \quad (8.2)$$

where $T \in \mathbb{R}_{\geq 0}$ plays the role of a temperature and $g_j^{(i)} = -\log[-\log U_j^{(i)}]$ with $U_j^{(i)} \sim \text{Uniform}(0, 1)$ samples of the uniform distribution. The $g_j^{(i)}$ are then samples of the so called Gumbel distribution and the resulting entries of the mask $m_j^{(i)} \sim \text{Concrete}(\phi_j^{(i)}, T)$ become samples of the concrete distribution. To better grasp the general properties of the weights, let us consider the extreme case in which the temperature goes to zero $T \rightarrow 0$. The softmax computation in Eq. (8.2) then smoothly approaches the discrete **argmax** computation for each row i (index is omitted):

$$\begin{aligned} \text{Concrete}(\phi_j, T \rightarrow 0) &= \mathbf{onehot}[\text{argmax}(\log \phi_j - g_j)] \in \mathbb{R}^d, \quad \text{with} \\ \text{Prob}\{m_j = 1\} &= \frac{\phi_j}{\sum_{l=1}^d \phi_l}, \end{aligned} \quad (8.3)$$

corresponding to the Gumbel-Max trick, a special instance of *reparametrization trick*. Sampling from the concrete distribution at zero temperature yields to an equivalent sampling from a categorical distribution with parameters $\phi^{(i)}$. Therefore, instead of a direct sampling of the mask, we rely on another distribution (the Uniform one at the lowest level) plus a deterministic function (sum of the logarithms of ϕ s and g s plus the **argmax** function). Eventually,

this is crucial for getting a differentiable map of the ϕ s. At this stage, the zero temperature appearing in the differentiable `Softmax` function makes it the non-differentiable `argmax` function but a finite value ensures $m^{(i)}(\phi, T)$ to be continuous relaxations of the `argmax` [136]. As long as the temperature is not zero, the parameters can be optimized through their gradients.

The plan for the automatic feature selection on a reconstruction task is hence the following: we choose as encoding function of the auto-encoder the layer specified in (8.1) called from now on a **concrete selector** layer. We initialize the matrix of the parameters $\phi \in \mathbb{R}^{d \times k}$ with small random positive numbers and the temperature T with a high value. The concrete selector layer is sampled from the concrete distribution but it is not sparse and linear combinations of the input features are stochastically explored. During the training, the temperature is lowered according to an *annealing schedule* that drives it to a very small value so that the filtering of the concrete selector becomes more and more sparse. Eventually, it selects only single input features in each row and the parameters ϕ and θ are optimized to obtain a good reconstruction.

Assuming that the procedure converges to a sufficiently low level of the reconstruction error, the autoencoder with the concrete selector is capable of reconstructing the entire input data of the training set from a subset of relevant features \mathcal{D} only. The other assumption is that, during training, the components of every row $\phi^{(i)}$ are converging to a sharp distributions peaked on one dominant component. When this is the case, the method succeeds in carrying out the automatic selection. For the sake of clarity, we sketch a pseudo-code of the training procedure in Algorithm 2.

While implementing the concrete selector and testing on the ES dataset, we actually ran into the problem of the same feature being selected multiple times. Instead, we want the concrete selector to be able to choose k input components which are different between each other. This problem is solved by considering only one vector of parameters $\phi \in \mathbb{R}^d$ whose components are interpreted as probabilities of the input features to be selected. Every row of the mask weights is then sampled directly on the same vector in a way that ensures M_ϕ to have orthonormal rows: once the temperature parameter is decreased to zero, the mask's rows are different one-hot vectors in the \mathbb{R}^d space.

Algorithm 2 Training of Concrete Autoencoder

-
- 1: Initialize concrete selector encoder with small positive ϕ
 - 2: Initialize decoder network g_θ with random initial θ
 - 3: Initialize initial temperature T_0 and final temperature $T_f = T_B$
 - 4: Initialize learning rate λ
 - 5: **for** epoch b of total B **do**
 - 6: update temperature $T = T_0 \left(\frac{T_B}{T_0}\right)^{-b/B}$
 - 7: **for** step c of total C **do**
 - 8: sample $m^{(i)} \sim \text{Concrete}(\phi^{(i)}, T) \quad \forall i \in [1, \dots, k]$
 - 9: reconstruct $\vec{x}_j = g_\theta[M_\phi \cdot \vec{x}_j] \quad \forall j$ in the batch
 - 10: compute the loss $\mathcal{L} = \text{mean}(\|\vec{x}_j - \hat{\vec{x}}_j\|^2)$
 - 11: optimize the parameters through gradients:

$$\theta \rightarrow \theta - \lambda \nabla_\theta \mathcal{L}$$

$$\alpha \rightarrow \alpha - \lambda \nabla_\alpha \mathcal{L}$$

12: **end for**

13: **end for**

8.1 Training for the ES dataset

We report here some details about the training of the CAE on the ES dataset for the BH model with $L = 256$ sites. The training is performed with the Adam optimizer [55] with a learning rate of 5×10^{-3} over a training set composed of around 500 examples from the interval $0 \leq U/t \leq 2.5$. We also considered again a validation set of examples from the $2.5 \leq U/t \leq 3$ interval. The concrete selector is initialized with random small and positive ϕ and the initial temperature is set to $T = 10$. The temperature is decreased afterwards to 0.01 in 150 epochs with an exponential falloff. The concrete selector is used as unique encoding layer of the auto-encoder: in this way we set a quite restrictive constraint affecting the maximum complexity of the encoding network in favour of interpretability. The architecture can be improved by employing the concrete selector as first layer and some additional encoding layers so that the interpretability is possible even though it is not referred to the bottleneck's features. In our implementation, the latent space is indeed forced to learn a subset of the input giving a fully explainable picture in terms of sectors of the ES. The decoding layer is made of three dense layers of dimension $d/2$, $d/2$ and d (the output one) and hyperbolic tangent activation functions except for the last layer that has Softmax in order to ensure that the output is a probability distribution. Dropout [57] with a probability of 0.3 was added to

prevent overfitting.




Figure 8.1: Evolution of the probability of the first k parameters ϕ (left) and of the trace of $(M_\phi)^T \cdot M_\phi$ (right) as a function of the number of iterations. The dimension of the latent space of the CAE is $k = 3$. The temperature is decreased exponentially from $T = 10$ to $T = 0.01$ and represented with a red line referred to the right-most axis in logarithmic scale.

During training, some metrics are monitored: for a given k , namely the cardinality of the subset of input features to be chosen by the concrete selector, we monitor the sum of the probabilities $p_{topk} = \sum_j^k p(\phi_j)$ of the first k parameters ϕ_j in magnitude. In a successful training, after an initial exploration of the possible combinations of input components, the CAE must converge to a situation where the biggest k parameters have a probability that sums up to one while all the others are suppressed as the temperature decreases. This is a good indication that the CAE has chosen the subset. Moreover, we also monitor the orthogonality of the rows of M_ϕ via the trace of $(M_\phi)^T \cdot M_\phi$ which must converge to k for selecting different input components. Clearly, the reconstruction error for both training and validation sets is also considered.

We report in Fig. 8.1 the evolution of the aforementioned metrics as a function of the number of iterations, for a CAE with a $k = 3$ latent space. The probability p_{topk}/k reaches approximately 1 after 20 epochs. Concerning the measure of orthogonality, the trace divided by k approaches one around the 80-th epoch where the temperature is of the order of 0.1. The reconstruction error and the validation error are comparable to each other and reach values around 0.015.

The evolution of the parameters ϕ normalized as probability values is reported




Figure 8.2: The probability (color scale) of the ES sectors for the 256 sites BH system to be selected by the concrete selector layer as the training proceeds. The sectors are labeled by symmetry number δN (first digit) and k ordering index (second digit) in the same notation described in section 6.3. After the random initialization with almost uniform probability, the layer becomes more and more sparse: the sector of the eigenvalue ξ_0^1 dominates over the others after only few iterations followed by ξ_0^0 and $\xi_{\pm 1}^1$. The temperature is represented in the bottom panel.

in Fig. 8.2. Every ϕ_j is uniquely associated to a specific component of the input, hence eigenvalue of the ES. When the p_{topk}/k reaches 1, the vector of parameters is already sparse meaning that the concrete selector has chosen the most relevant features for the reconstruction. Since from this point the reconstruction error keeps decreasing, we suppose that after the first exploration, once the features are selected, the optimization is mainly done on the level of the decoder, in order to enhance the reconstruction of the entire ES.

8.2 Results for 256 sites BH model




Figure 8.3: Original eigenvalues (left) and reconstructed eigenvalues by the first CAE model with $k = 3$ (right) for the 256 sites BH model as a function of U/t . The CAE was trained with configurations from the region $0 \leq U/t \leq 2.5$. Despite the low dimensionality of the latent space, the CAE is able to reproduce quite faithfully the ES even outside the training window, up to the transition region.

We present here the results of three different CAEs on the ES dataset for the BH model with $L = 256$ sites. The first model has a latent dimension of $k = 3$ and after the training the selected eigenvalues are $\xi_{\delta N=0}^{k=0}$, ξ_0^1 and ξ_{-1}^1 , where for the notation about the index of the eigenvalues we refer to the same as in section 6.3. The training is performed with configurations from the region $0 \leq U/t \leq 2.5$. It is very interesting to examine in depth the strategy that the CAE chooses for the sake of reconstruction now that we have precise information about the eigenvalues that it takes into account. For this CAE the first one is the lowest, that sets the scale of the ES and the position of the minimum for the first conformal parabola of the ES. As we have seen in section 6.3, the second eigenvalue is fundamental for computing the ratio that allows to collapse all the equidistant parabolas to horizontal lines (see Eq. (6.3)). A last, the third eigenvalue is fundamental for computing the η quantity (see section 6.4 and [98]) necessary for determining the critical point in the spirit of level spectroscopy. It is quite surprising that the CAE is able to learn automatically the best strategy to reconstruct the parabolic shapes which are the main characteristics of the gapless regime without having any notion about them. On the other hand, it demonstrates the promising potential of the CAE regarding the explainability of the machine learning model for distinguishing between phases of the system. With those three eigenvalues the CAE is able to properly reconstruct the ES for all the values of the interaction strength U/t up to the region of the transition to the gapped




Figure 8.4: Reconstruction of the symmetry resolved ES for the 256 sites BH model with $k = 3$ CAE for three values of the interaction strength: on the left for $U/t = 0.5$ deep in the SF phase, in the central panel for $U/t = 3$ still in the SF regime but closer to the transition and outside the training window and on the right for $U/t = 4$ in the gapped phase. The diamonds represent the reconstructed eigenvalues while the dots are the ground truth. The sectors chosen by the concrete selector are highlighted with the blue rectangles.

phase without having any notion of the U value at disposal. The reconstruction of the ES is reported in Fig. 8.3 and compared with the original spectrum. Moreover, we show in Fig. 8.4 the comparison between the original ES and the reconstructed one with a symmetry resolved representation for three values of the interaction strength: on the left for $U/t = 0.5$ deep in the SF phase where the training is done, in the central panel for $U/t = 3$ still in the SF regime but closer to the transition and outside the training region. On the right we show the comparison for $U/t = 4$ in the gapped phase. The chosen sectors are highlighted with blue rectangles around the true eigenvalues (dots), the reconstructed ones are represented with diamonds. Even in the gapped phase, the CAE keeps building parabolic ES, therefore failing in the reproduction task and obtaining a reconstruction loss. We observed that with only 3 available degrees of freedom in the latent space, the generalization abilities of the CAE are quite poor and the shapes of the parabolas are almost always the same for all the values of the coupling.

In Fig. 8.5 we present the analogue plot of Fig. 8.4 for another instance of the CAE with latent dimension $k = 3$. Due to the random initialization, the selected eigenvalues can be sometimes different and for this instance they correspond to ξ_0^1 , ξ_1^1 and ξ_{-1}^1 . Those are the lowest sectors of the second

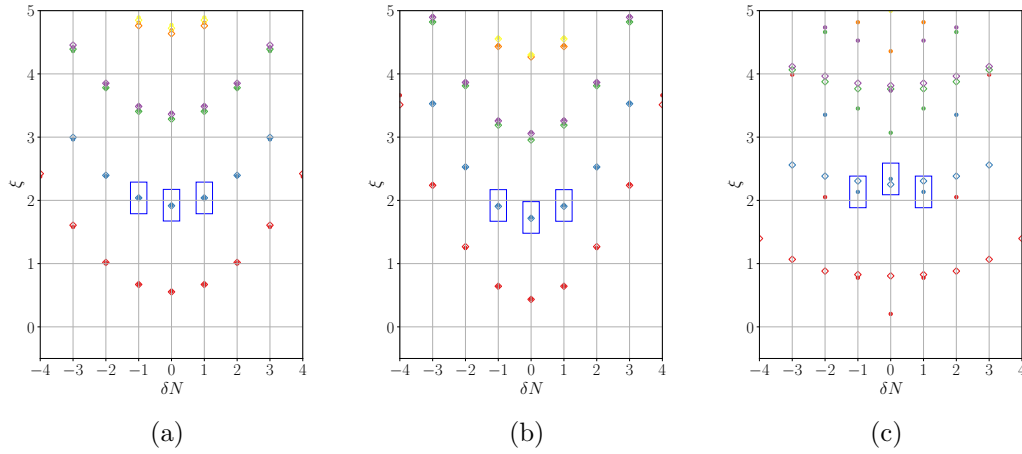


Figure 8.5: Reconstruction of the symmetry resolved ES for the 256 sites BH model with $k = 3$ CAE for three values of the interaction strength: $U/t = 0.5$ (left), $U/t = 3$ (center), $U/t = 4$ (right). The diamonds represent the reconstructed eigenvalues while the points are the original ones. The chosen sectors by the concrete selector are highlighted with the blue rectangles. For this model they are different from the previous one and the reconstruction quality in the gapped phase is affected by the specific selection.

parabola, meaning that the CAE is again focused on the parabolic envelopes and has probably encoded the information about the vertical shift between the towers somehow in the decoding layers. As a consequence: in the MI phase the CAE reconstruct parabolas based on the relative positions of the observed sectors without success. Indeed, in the gapped phase the ξ_0^1 eigenvalue progressively drifts away from ξ_0^0 , destroying the equidistant parabolas and reaching a peculiar U and W shape for the first and second order of eigenvalues respectively, deep in the MI. In this phase, the CAE reconstruction results in flat structures.

The last example is a CAE with latent dimension $k = 4$ with a training region further from the transition $0 \leq U/t \leq 1.5$. The chosen eigenvalues are found to be almost always ξ_0^0 , ξ_0^1 , ξ_{-1}^1 and ξ_1^1 even with different random initial weights. The reconstruction of the ES is shown in Fig. 8.6. After the transition has occurred, the CAE always outputs equidistant parabolas which are never flat as in the previous case. Also for this configuration we plot in Fig. 8.7 the spectra for $U/t = 0.5$, $U/t = 3$ and $U/t = 4$. Looking at the loss profile as a function of the coupling U/t , already the $k = 4$ CAE is able to provide very good results for the anomaly detection. The very simple concrete selector is sufficient to automatically grasp the underlying features in the ES, which




Figure 8.6: Reconstructed eigenvalues by the third CAE model with $k = 4$ for the 256 sites BH model as a function of U/t (for the original spectrum see Fig. 8.3). The CAE was trained with configurations from the region $0 \leq U/t \leq 1.5$ so much smaller than in the previous cases.

are very structured in the gapless phase and suggests to look at the most reasonable eigenvalues to derive all the others.




Figure 8.7: Reconstruction of the symmetry resolved ES for the 256 sites BH model with $k = 3$ CAE for three values of the interaction strength: $U/t = 0.5$ (left), $U/t = 3$ (center), $U/t = 4$ (right). The diamonds represent the reconstructed eigenvalues while the points are the original ones. The chosen sectors by the concrete selector are highlighted with the blue rectangles and are concretely related to the most important low-lying eigenvalues of the ES for the gapless phase. After the transition has occurred at $U/t = 4$, the CAE always reconstructs equidistant parabolas.

8.3 Conclusions

We have shown that the CAE represents a valuable tool for the sake of pursuing an automatic feature selection of the input features while performing an anomaly detection protocol to detect phase transitions. We employed the CAE to learn how to reconstruct the ES of the 256 sites BH model by automatically selecting the most important eigenvalues of the ES for the gapless phase and reconstructing all the others from them. The advantages of the CAE for the feature selection are mainly its generality and ease of use. We believe that the inclusion of the CAE in the GAN architecture studied in chapter 6 can be a very promising approach since it offers the added-value of interpretability. We aim at evaluating the performance of the CAE for other datasets and systems in the future. It can be of particular use in simplifying analyses involving very raw data with a large number of related quantities, especially when the characterization of the phases relying on such quantities chosen for the detection is not clear.

Chapter 9

Discussion and Conclusions

In this thesis, we provided a data-driven approach to the detection of quantum phase transitions. The goal of this study was twofold: to identify and to benchmark a machine learning protocol as a tool for discerning among different quantum phases with minimal prior knowledge about the system under investigation and, on the other hand, to explore some general quantities to be used as features for the machine learning model. Moreover, we wanted to address the problem of interpretability, which is an hot-topic in the artificial intelligence community, and aims at making the prediction of the model more transparent to the user.

Firstly, we propose the Entanglement Spectrum (ES) as a reliable quantity for machine-driven detection of the elusive Berezinskii-Kosterlitz-Thouless (BKT) transition in one-dimensional quantum systems. By utilizing a Generative Adversarial Network (GAN) trained on the spectra belonging to the gapless phase, we develop an anomaly detection scheme to effectively discern the changes in ES patterns when feeding the trained network with data from other phases. Applying the scheme to numerical data simulated with MPSs from various models including the XXZ spin chain, the Bose Hubbard model, and the two-component Bose Hubbard model, we obtain promising results and estimate critical points even from small system sizes' data. While our method is tested on ES, it can be applied to any representative quantity of the many-body state embedding some changes along the phase transitions, automatically providing a map of the phase diagram.

Secondly, we demonstrate the power of the probability density function (PDF)

of the bosonic occupation number in exploring the phase diagram of quantum many-body systems. The PDF, along with simple fits, serves as an accessible and informative intermediate quantity between the full ES and the number fluctuations. Our results for the extended Bose-Hubbard model align well with sophisticated traditional and machine learning analyses. We believe that the PDF combines the advantages of experimental accessibility and rich information about the full many-body state, making it suitable for automatic detection protocols. We also suggest the extension of this method to non-zero temperature regimes and systems beyond one dimension.

Lastly, we showcase the value of the Concrete Autoencoder (CAE) as a tool for feature selection in anomaly detection protocols for phase transitions. By using the CAE, we successfully reconstruct the ES of the 256-site Bose-Hubbard model, automatically selecting the most significant eigenvalues for the gapless phase and reproducing the remaining ones. The CAE's generality and ease of use make it a promising addition to the GAN architecture explored for the anomaly detection scheme, providing interpretability alongside its functionality. Future work involves evaluating the CAE's performance on different datasets and systems, particularly for raw data with numerous related quantities and unclear phase characterizations.

Appendix A

Luttinger parameters

The low-energy behaviour of all the models considered in this thesis in the gapless phase can be effectively described by the Tomonaga-Luttinger liquid theory [77]. The latter theory is actually able to characterize the properties of any one-dimensional massless system by means of two phenomenological parameters, namely the Luttinger parameters K and u , appearing in the hamiltonian:

$$H_{LL} = \frac{u}{2\pi} \int \left[K(\partial_x \theta)^2 + \frac{1}{K}(\partial_x \phi)^2 \right] dx \quad (\text{A1})$$

where the field ϕ encodes the long wave-length fluctuations around the average density of particles $\rho(x) = \rho_0 + \frac{1}{\pi}(\partial_x \phi)$ and $(\partial_x \theta) = \Pi(x)$ is the canonically conjugate momentum of ϕ . The two Luttinger parameters are computed by fitting the low-energy spectrum of the system and once they are fixed, all the properties of the system are determined. Moreover, all the asymptotic properties of the correlation functions can be exactly obtained as a function of those parameters within the theory. Clearly, due to such a universal effective description, also the bipartite entanglement properties are common to all the models responding to the theory. In particular, the entanglement entropy leading scaling for the infinite system is the one of a free bosonic CFT with a central charge of $c = 1$ and the fluctuations exhibit the same specific scaling up to a prefactor depending on K as extensively described in [123].

In the following we present the results of the fitting of the Luttinger parameters for the BH model and the two-component BH model. In the first case the value was extrapolated from the density fluctuations of a subsystem as the second moment of the PDF described in chapter 7. Thanks to the predicted values

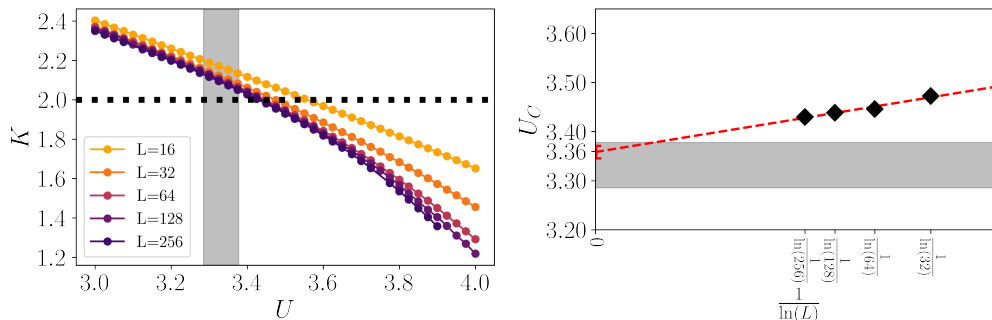


Figure A1: Detection of the BKT transition by means of the second moment of the PDF distribution, i.e. the number fluctuations (see [124]). (*left*) the curves of the effective Luttinger parameter K^* values for different sizes of the system as a function of U ($t = 1$ and $V = 0$). The critical true value is expected at $K = 2$ from the Luttinger theory (dotted black line). The gray regions stand for the single standard deviation confidence interval as from the estimates of literature summarized in Table 1 of [107]. (*right*) shows with diamonds the pseudo-critical values extracted from the crossing of the K^* curves with $K = 2$, as a function of the reciprocal of the log-size. With a linear fit (red dashed line) the extrapolation of U_C to the thermodynamic limit is the intercept.

of K at the transition point between the SF phase and the MI phase, we were able to compute the critical value U_C of the interaction strength at which the BKT transition occurs. In particular, the determination of the critical point is done resorting to the expected Luttinger parameter $K = 2$. The procedure is also experimentally realizable and is normally carried out via a fit of the fluctuations trend versus the (log of) the bipartition length of the subsystem at fixed system size. In order to obtain the true value of K , one must perform an extrapolation to the thermodynamic limit, therefore considering different system's sizes.

We report the results of such analysis in Figure A1. On the left, the effective K^* per size of the system is shown in the region of the BKT. We plotted as a guide-to-the-eye a gray shaded area representing the critical region as from the numerical studies listed in [107] and as gray dashed line the expected critical value of the Luttinger parameter $K = 2$. We derive the values of the pseudo-critical U where the effective K^* crosses with $K = 2$ and we perform a fit against the inverse log of the system's size (see right plot of A1). Eventually we obtain the estimate of the critical coupling $U_C = 3.36 \pm 0.01$ which is in perfect agreement with the results in the literature.

For the BH2S model, the Luttinger description in the 2SF phase for our system




Figure A2: The curves of the effective Luttinger parameter K_S^* values for different sizes of the system as a function of U_{AB}/U ($U = 10$). The gray regions stands for the expected interval for the transition as predicted in Fig. 6.12.

corresponds to the Hamiltonian of two coupled Luttinger liquids. The Hamiltonian can be diagonalised by introducing the density (D) and spin/polarisation (S) channels:

$$H_{\mu=D/S} = \frac{u_\mu}{2\pi} \int \left[K_\mu (\partial_x \phi_\mu)^2 + \frac{1}{K_\mu} (\partial_x \theta_\mu)^2 \right] dx, \quad (\text{A2})$$

where $\phi_{D(S)} = (\phi_A \pm \phi_B)/\sqrt{2}$ and $\theta_{D(S)} = (\theta_A \pm \theta_B)/\sqrt{2}$ are the bosonic fields related to the fluctuations of the phase and the amplitude of the total density (spin) of the two coupled superfluids. There's an additional non-linear coupling between the densities of the two Luttinger liquids due to the inter-species interactions which can be perturbatively accounted for by a term proportional to $U_{AB} \cos(2\sqrt{2}\theta_S)$. This term is irrelevant in the 2SF phase and relevant in the PSF phase. As long as the Hamiltonian Eq. (A2) holds, an algebraic decay characterises the correlation functions (a.k.a. quasi-long-range order) of Eq. 5.5 with a precise dependence on the Luttinger parameters K_D and K_S :

$$\begin{aligned} G_\alpha(x) &\propto |d|^{-\frac{1}{4K_D} - \frac{1}{4K_S}}, \\ R_D(x) &\propto |d|^{-\frac{1}{K_D}}, \\ R_S(x) &\propto |d|^{-\frac{1}{K_S}}. \end{aligned} \quad (\text{A3})$$

Here we expressed the algebraic decay in terms of the natural measure of the distances between sites on a ring geometry, i.e., the *chord* function [123]:

$$d(x/L) = \frac{L}{\pi} \sin\left(\frac{\pi x}{L}\right), \quad (\text{A4})$$

where L is the number of sites and $x \in \mathbb{N}$ the linear distance between the sites. For very large rings the expression further simplifies according to the substitution $d \rightarrow x$. Since we are interested in the region between the 2SF phase and the PSF phase, we show in Figure A2 the results of the fit of the Luttinger parameter K_S which is expected to jump to zero at the critical U_{AB} for which the system enters the PSF phase. Since our systems are finite, K_S^* exists in both the phases but the error in the fit becomes higher and higher after the critical coupling.

Bibliography

- [1] A. Einstein, B. Podolsky, and N. Rosen. “Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?” In: *Phys. Rev.* 47 (10 May 1935), pp. 777–780. DOI: [10.1103/PhysRev.47.777](https://doi.org/10.1103/PhysRev.47.777) (cit. on p. 4).
- [2] J. S. Bell. “On the Einstein Podolsky Rosen paradox”. In: *Physics Physique Fizika* 1 (3 Nov. 1964), pp. 195–200. DOI: [10.1103/PhysicsPhysiqueFizika.1.195](https://doi.org/10.1103/PhysicsPhysiqueFizika.1.195) (cit. on pp. 4, 5).
- [3] Alain Aspect, Jean Dalibard, and Gérard Roger. “Experimental Test of Bell’s Inequalities Using Time-Varying Analyzers”. In: *Phys. Rev. Lett.* 49 (25 Dec. 1982), pp. 1804–1807. DOI: [10.1103/PhysRevLett.49.1804](https://doi.org/10.1103/PhysRevLett.49.1804) (cit. on p. 5).
- [4] Gregor Weihs et al. “Violation of Bell’s Inequality under Strict Einstein Locality Conditions”. In: *Phys. Rev. Lett.* 81 (23 Dec. 1998), pp. 5039–5043. DOI: [10.1103/PhysRevLett.81.5039](https://doi.org/10.1103/PhysRevLett.81.5039) (cit. on p. 5).
- [5] Alain Aspect. “Bell’s inequality test: more ideal than ever”. In: *Nature* 398.6724 (Mar. 1999), pp. 189–190. DOI: [10.1038/18296](https://doi.org/10.1038/18296) (cit. on p. 5).
- [6] Jian-Wei Pan et al. “Experimental test of quantum nonlocality in three-photon Greenberger–Horne–Zeilinger entanglement”. In: *Nature* 403.6769 (Feb. 2000), pp. 515–519. DOI: [10.1038/35000514](https://doi.org/10.1038/35000514) (cit. on p. 5).
- [7] J. von Neumann. “Wahrscheinlichkeitstheoretischer Aufbau der Quantenmechanik”. In: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1927 (1927), pp. 245–272 (cit. on p. 6).
- [8] C. E. Shannon. “A mathematical theory of communication”. In: *The Bell System Technical Journal* 27.3 (1948), pp. 379–423. DOI: [10.1002/j.1538-7305.1948.tb01338.x](https://doi.org/10.1002/j.1538-7305.1948.tb01338.x) (cit. on p. 6).
- [9] Charles H. Bennett et al. “Concentrating partial entanglement by local operations”. In: *Phys. Rev. A* 53 (4 Apr. 1996), pp. 2046–2052. DOI: [10.1103/PhysRevA.53.2046](https://doi.org/10.1103/PhysRevA.53.2046) (cit. on p. 6).
- [10] Martin B. Plenio and Shashank Virmani. “An Introduction to entanglement measures”. In: *Quant. Inf. Comput.* 7 (2007), pp. 1–51. arXiv: [quant-ph/0504163](https://arxiv.org/abs/quant-ph/0504163) (cit. on p. 6).

- [11] Luigi Amico et al. “Entanglement in many-body systems”. In: *Rev. Mod. Phys.* 80 (2 May 2008), pp. 517–576. DOI: [10.1103/RevModPhys.80.517](https://doi.org/10.1103/RevModPhys.80.517) (cit. on pp. 6, 7).
- [12] G. Vidal et al. “Entanglement in Quantum Critical Phenomena”. In: *Phys. Rev. Lett.* 90 (22 June 2003), p. 227902. DOI: [10.1103/PhysRevLett.90.227902](https://doi.org/10.1103/PhysRevLett.90.227902) (cit. on pp. 8, 9, 89).
- [13] Jacob D. Bekenstein. “Black Holes and Entropy”. In: *Phys. Rev. D* 7 (8 Apr. 1973), pp. 2333–2346. DOI: [10.1103/PhysRevD.7.2333](https://doi.org/10.1103/PhysRevD.7.2333) (cit. on p. 8).
- [14] S. W. Hawking. “Black hole explosions?” In: *Nature* 248.5443 (Mar. 1974), pp. 30–31. DOI: [10.1038/248030a0](https://doi.org/10.1038/248030a0) (cit. on p. 8).
- [15] Mark Srednicki. “Entropy and area”. In: *Phys. Rev. Lett.* 71 (5 Aug. 1993), pp. 666–669. DOI: [10.1103/PhysRevLett.71.666](https://doi.org/10.1103/PhysRevLett.71.666) (cit. on p. 8).
- [16] Pasquale Calabrese and John Cardy. “Entanglement entropy and quantum field theory”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2004.06 (June 2004), P06002. DOI: [10.1088/1742-5468/2004/06/P06002](https://doi.org/10.1088/1742-5468/2004/06/P06002) (cit. on pp. 8, 9, 89).
- [17] M B Hastings. “An area law for one-dimensional quantum systems”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2007.08 (Aug. 2007), P08024. DOI: [10.1088/1742-5468/2007/08/P08024](https://doi.org/10.1088/1742-5468/2007/08/P08024) (cit. on p. 9).
- [18] Christoph Holzhey, Finn Larsen, and Frank Wilczek. “Geometric and renormalized entropy in conformal field theory”. In: *Nuclear Physics B* 424.3 (1994), pp. 443–467. DOI: [https://doi.org/10.1016/0550-3213\(94\)90402-2](https://doi.org/10.1016/0550-3213(94)90402-2) (cit. on p. 9).
- [19] X. G. Wen. “Vacuum degeneracy of chiral spin states in compactified space”. In: *Phys. Rev. B* 40 (10 Oct. 1989), pp. 7387–7390. DOI: [10.1103/PhysRevB.40.7387](https://doi.org/10.1103/PhysRevB.40.7387) (cit. on p. 9).
- [20] Alexei Kitaev and John Preskill. “Topological Entanglement Entropy”. In: *Phys. Rev. Lett.* 96 (11 Mar. 2006), p. 110404. DOI: [10.1103/PhysRevLett.96.110404](https://doi.org/10.1103/PhysRevLett.96.110404) (cit. on p. 9).
- [21] Joseph J. Bisognano and Eyvind H. Wichmann. “On the duality condition for quantum fields”. In: *Journal of Mathematical Physics* 17.3 (1976), pp. 303–321. DOI: [10.1063/1.522898](https://doi.org/10.1063/1.522898). eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.522898> (cit. on p. 12).
- [22] Pasquale Calabrese and Alexandre Lefevre. “Entanglement spectrum in one-dimensional systems”. In: *Phys. Rev. A* 78.3 (Sept. 2008). DOI: [10.1103/physreva.78.032329](https://doi.org/10.1103/physreva.78.032329) (cit. on pp. 12, 69).
- [23] Hui Li and F. D. M. Haldane. “Entanglement Spectrum as a Generalization of Entanglement Entropy: Identification of Topological Order in Non-Abelian Fractional Quantum Hall Effect States”. In: *Phys. Rev. Lett.* 101 (1 July 2008), p. 010504. DOI: [10.1103/PhysRevLett.101.010504](https://doi.org/10.1103/PhysRevLett.101.010504) (cit. on pp. 12, 70, 89).

- [24] Nicolas Lafflorencie. “Quantum entanglement in condensed matter systems”. In: *Physics Reports* 646 (Aug. 2016), pp. 1–59. DOI: [10.1016/j.physrep.2016.06.008](https://doi.org/10.1016/j.physrep.2016.06.008) (cit. on p. 12).
- [25] Kenneth G. Wilson. “Renormalization Group and Critical Phenomena”. In: *Phys. Rev. B* 4 (9 Nov. 1971), pp. 3174–3183. DOI: [10.1103/PhysRevB.4.3174](https://doi.org/10.1103/PhysRevB.4.3174) (cit. on p. 13).
- [26] Kenneth G. Wilson. “The renormalization group: Critical phenomena and the Kondo problem”. In: *Rev. Mod. Phys.* 47 (4 Oct. 1975), pp. 773–840. DOI: [10.1103/RevModPhys.47.773](https://doi.org/10.1103/RevModPhys.47.773) (cit. on p. 13).
- [27] S. R. White and R. M. Noack. “Real-space quantum renormalization groups”. In: *Phys. Rev. Lett.* 68 (24 June 1992), pp. 3487–3490. DOI: [10.1103/PhysRevLett.68.3487](https://doi.org/10.1103/PhysRevLett.68.3487) (cit. on pp. 13, 14).
- [28] Steven R. White. “Density matrix formulation for quantum renormalization groups”. In: *Phys. Rev. Lett.* 69 (19 Nov. 1992), pp. 2863–2866. DOI: [10.1103/PhysRevLett.69.2863](https://doi.org/10.1103/PhysRevLett.69.2863) (cit. on pp. 13–15).
- [29] U. Schollwöck. “The density-matrix renormalization group”. In: *Rev. Mod. Phys.* 77 (1 Apr. 2005), pp. 259–315. DOI: [10.1103/RevModPhys.77.259](https://doi.org/10.1103/RevModPhys.77.259) (cit. on p. 13).
- [30] Ian Affleck et al. “Rigorous results on valence-bond ground states in antiferromagnets”. In: *Phys. Rev. Lett.* 59 (7 Aug. 1987), pp. 799–802. DOI: [10.1103/PhysRevLett.59.799](https://doi.org/10.1103/PhysRevLett.59.799) (cit. on p. 17).
- [31] Stellan Östlund and Stefan Rommer. “Thermodynamic Limit of Density Matrix Renormalization”. In: *Phys. Rev. Lett.* 75 (19 Nov. 1995), pp. 3537–3540. DOI: [10.1103/PhysRevLett.75.3537](https://doi.org/10.1103/PhysRevLett.75.3537) (cit. on p. 17).
- [32] J. Dukelsky et al. “Equivalence of the variational matrix product method and the density matrix renormalization group applied to spin chains”. In: *Europhysics Letters* 43.4 (Aug. 1998), p. 457. DOI: [10.1209/epl/i1998-00381-x](https://doi.org/10.1209/epl/i1998-00381-x) (cit. on p. 17).
- [33] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. In: *Annals of Physics* 326.1 (2011). January 2011 Special Issue, pp. 96–192. DOI: <https://doi.org/10.1016/j.aop.2010.09.012> (cit. on pp. 17, 29).
- [34] Fabio Franchini. “An Introduction to Integrable Techniques for One-Dimensional Quantum Systems”. In: Berlin, Heidelberg: Springer International Publishing, 2017. DOI: [10.1007/978-3-319-48487-7](https://doi.org/10.1007/978-3-319-48487-7) (cit. on pp. 17, 57, 71).
- [35] David Landau and Kurt Binder. *A Guide to Monte Carlo Simulations in Statistical Physics*. USA: Cambridge University Press, 2005 (cit. on p. 17).
- [36] F. Verstraete and J. I. Cirac. “Renormalization algorithms for quantum-many body systems in two and higher dimensions”. In: (July 2004). arXiv: [cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066) (cit. on p. 17).

- [37] Pietro Silvi et al. “The Tensor Networks Anthology: Simulation techniques for many-body quantum lattice systems”. In: *SciPost Physics Lecture Notes* (Mar. 2019). DOI: [10.21468/scipostphyslectnotes.8](https://doi.org/10.21468/scipostphyslectnotes.8) (cit. on pp. 18, 27, 31, 33).
- [38] Gregory M. Crosswhite and Dave Bacon. “Finite automata for caching in matrix product algorithms”. In: *Phys. Rev. A* 78.1 (July 2008). DOI: [10.1103/physreva.78.012356](https://doi.org/10.1103/physreva.78.012356) (cit. on p. 25).
- [39] Michael P. Zaletel et al. “Time-evolving a matrix product state with long-ranged interactions”. In: *Phys. Rev. B* 91.16 (Apr. 2015). DOI: [10.1103/physrevb.91.165112](https://doi.org/10.1103/physrevb.91.165112) (cit. on p. 25).
- [40] Garnet Kin-Lic Chan et al. *Matrix Product Operators, Matrix Product States, and ab initio Density Matrix Renormalization Group algorithms*. 2016. DOI: [10.48550/ARXIV.1605.02611](https://doi.org/10.48550/ARXIV.1605.02611). eprint: [1605.02611](https://arxiv.org/abs/1605.02611) (cit. on p. 25).
- [41] S. Kung and D. Lin. “Optimal Hankel-norm model reductions: Multi-variable systems”. In: *IEEE Transactions on Automatic Control* 26.4 (1981), pp. 832–852. DOI: [10.1109/TAC.1981.1102736](https://doi.org/10.1109/TAC.1981.1102736) (cit. on p. 26).
- [42] Román Orús. “A practical introduction to tensor networks: Matrix product states and projected entangled pair states”. In: *Annals of Physics* 349 (Oct. 2014), pp. 117–158. DOI: [10.1016/j.aop.2014.06.013](https://doi.org/10.1016/j.aop.2014.06.013) (cit. on p. 27).
- [43] Jane K. Cullum and Ralph A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations*. Society for Industrial and Applied Mathematics, 2002. DOI: [10.1137/1.9780898719192](https://doi.org/10.1137/1.9780898719192). eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9780898719192> (cit. on p. 28).
- [44] Gerard L. G. Sleijpen et al. “Jacobi-davidson type methods for generalized eigenproblems and polynomial eigenproblems”. In: *BIT Numerical Mathematics* 36.3 (Sept. 1996), pp. 595–633. DOI: [10.1007/BF01731936](https://doi.org/10.1007/BF01731936) (cit. on p. 28).
- [45] Peter Pippian, Steven R. White, and Hans Gerd Evertz. “Efficient matrix-product state method for periodic boundary conditions”. In: *Phys. Rev. B* 81 (8 Feb. 2010), p. 081103. DOI: [10.1103/PhysRevB.81.081103](https://doi.org/10.1103/PhysRevB.81.081103) (cit. on p. 33).
- [46] Andreas Haller, Matteo Rizzi, and Michele Filippone. “Drude weight increase by orbital and repulsive interactions in fermionic ladders”. In: *Phys. Rev. Res.* 2 (2 Apr. 2020), p. 023058. DOI: [10.1103/PhysRevResearch.2.023058](https://doi.org/10.1103/PhysRevResearch.2.023058) (cit. on p. 33).
- [47] Junyi Chai et al. “Deep learning in computer vision: A critical review of emerging techniques and application scenarios”. In: *Machine Learning with Applications* 6 (2021), p. 100134. DOI: <https://doi.org/10.1016/j.mlwa.2021.100134> (cit. on p. 35).

- [48] Dong Yu and Li Deng. *Automatic Speech Recognition: A Deep Learning Approach*. Springer Publishing Company, Incorporated, 2014 (cit. on p. 35).
- [49] Albert Gatt and Emiel Kraahmer. “Survey of the State of the Art in Natural Language Generation: Core Tasks, Applications and Evaluation”. In: *J. Artif. Int. Res.* 61.1 (Jan. 2018), pp. 65–170 (cit. on p. 35).
- [50] Alankrita Aggarwal, Mamta Mittal, and Gopi Battineni. “Generative adversarial network: An overview of theory and applications”. In: *International Journal of Information Management Data Insights* 1.1 (2021), p. 100004. DOI: <https://doi.org/10.1016/j.jjime.2020.100004> (cit. on pp. 35, 46, 67).
- [51] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book, 2018 (cit. on p. 35).
- [52] Michael A. Nielsen. *Neural Networks and Deep Learning*. <http://neuralnetworksanddeeplearning.com/>. Determination Press, 2015 (cit. on p. 35).
- [53] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016 (cit. on pp. 35, 39, 44, 67).
- [54] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. USA: Oxford University Press, Inc., 1995 (cit. on pp. 35, 69).
- [55] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2014. DOI: [10.48550/ARXIV.1412.6980](https://doi.org/10.48550/ARXIV.1412.6980). eprint: [1412.6980](https://arxiv.org/abs/1412.6980) (cit. on pp. 39, 104).
- [56] Jan Kukačka, Vladimir Golkov, and Daniel Cremers. *Regularization for Deep Learning: A Taxonomy*. 2017. arXiv: [1710.10686](https://arxiv.org/abs/1710.10686) [cs.LG] (cit. on p. 43).
- [57] Nitish Srivastava et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958 (cit. on pp. 43, 104).
- [58] Hervé Abdi and Lynne J. Williams. “Principal component analysis”. In: *WIREs Computational Statistics* 2.4 (2010), pp. 433–459. DOI: <https://doi.org/10.1002/wics.101>. eprint: <https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wics.101> (cit. on pp. 45, 76).
- [59] Subir Sachdev. *Quantum Phase Transitions*. 2nd ed. Cambridge University Press, 2011. DOI: [10.1017/CBO9780511973765](https://doi.org/10.1017/CBO9780511973765) (cit. on p. 49).
- [60] Amit Dutta et al. *Quantum Phase Transitions in Transverse Field Spin Models: From Statistical Physics to Quantum Information*. Cambridge University Press, 2015. DOI: [10.1017/CBO9781107706057](https://doi.org/10.1017/CBO9781107706057) (cit. on p. 49).
- [61] Daniele Contessi et al. “Detection of Berezinskii-Kosterlitz-Thouless transition via Generative Adversarial Networks”. In: *SciPost Phys.* 12 (2022), p. 107. DOI: [10.21468/SciPostPhys.12.3.107](https://doi.org/10.21468/SciPostPhys.12.3.107) (cit. on pp. 55, 66).

- [62] Daniele Contessi, Alessio Recati, and Matteo Rizzi. “Phase diagram detection via Gaussian fitting of number probability distribution”. In: *Phys. Rev. B* 107 (12 Mar. 2023), p. L121403. DOI: [10.1103/PhysRevB.107.L121403](https://doi.org/10.1103/PhysRevB.107.L121403) (cit. on p. 55).
- [63] Daniele Contessi et al. “Collisionless drag for a one-dimensional two-component Bose-Hubbard model”. In: *Phys. Rev. Research* 3 (2 May 2021), p. L022017. DOI: [10.1103/PhysRevResearch.3.L022017](https://doi.org/10.1103/PhysRevResearch.3.L022017) (cit. on pp. 55, 62, 63, 80).
- [64] Giuseppe Carleo et al. “Machine learning and the physical sciences”. In: *Rev. Mod. Phys.* 91 (4 Dec. 2019), p. 045002. DOI: [10.1103/RevModPhys.91.045002](https://doi.org/10.1103/RevModPhys.91.045002) (cit. on p. 55).
- [65] Juan Carrasquilla. “Machine learning for quantum matter”. In: *Advances in Physics: X* 5.1 (2020), p. 1797528. DOI: [10.1080/23746149.2020.1797528](https://doi.org/10.1080/23746149.2020.1797528). eprint: <https://doi.org/10.1080/23746149.2020.1797528> (cit. on p. 55).
- [66] Lei Wang. “Discovering phase transitions with unsupervised learning”. In: *Phys. Rev. B* 94 (19 Nov. 2016), p. 195105. DOI: [10.1103/PhysRevB.94.195105](https://doi.org/10.1103/PhysRevB.94.195105) (cit. on p. 55).
- [67] Juan Carrasquilla and Roger G. Melko. “Machine learning phases of matter”. In: *Nature Physics* 13.5 (Feb. 2017), pp. 431–434. DOI: [10.1038/nphys4035](https://doi.org/10.1038/nphys4035) (cit. on p. 55).
- [68] Evert P.L. Van Nieuwenburg, Ye Hua Liu, and Sebastian D. Huber. “Learning phase transitions by confusion”. In: *Nature Physics* 13.5 (2017), pp. 435–439. DOI: [10.1038/nphys4037](https://doi.org/10.1038/nphys4037) (cit. on pp. 55, 70, 89).
- [69] Sebastian J. Wetzel. “Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders”. In: *Phys. Rev. E* 96 (2 Aug. 2017), p. 022140. DOI: [10.1103/PhysRevE.96.022140](https://doi.org/10.1103/PhysRevE.96.022140) (cit. on p. 55).
- [70] Wenjian Hu, Rajiv R. P. Singh, and Richard T. Scalettar. “Discovering phases, phase transitions, and crossovers through unsupervised machine learning: A critical examination”. In: *Phys. Rev. E* 95 (6 June 2017), p. 062122. DOI: [10.1103/PhysRevE.95.062122](https://doi.org/10.1103/PhysRevE.95.062122) (cit. on p. 55).
- [71] Maciej Koch-Janusz and Zohar Ringel. “Mutual information, neural networks and the renormalization group”. In: *Nature Physics* 14.6 (Mar. 2018), pp. 578–582. DOI: [10.1038/s41567-018-0081-4](https://doi.org/10.1038/s41567-018-0081-4) (cit. on p. 55).
- [72] Patrick Huembeli et al. “Automated discovery of characteristic features of phase transitions in many-body localization”. In: *Phys. Rev. B* 99 (10 Mar. 2019), p. 104106. DOI: [10.1103/PhysRevB.99.104106](https://doi.org/10.1103/PhysRevB.99.104106) (cit. on p. 55).
- [73] Korbinian Kottmann et al. “Unsupervised Phase Discovery with Deep Anomaly Detection”. In: *Phys. Rev. Lett.* 125.17 (2020), pp. 1–7. DOI: [10.1103/PhysRevLett.125.170603](https://doi.org/10.1103/PhysRevLett.125.170603) (cit. on pp. 55, 59, 66, 70, 78, 89, 92, 99).

- [74] T. Mendes-Santos et al. “Unsupervised Learning Universal Critical Behavior via the Intrinsic Dimension”. In: *Phys. Rev. X* 11 (1 Feb. 2021), p. 011040. DOI: [10.1103/PhysRevX.11.011040](https://doi.org/10.1103/PhysRevX.11.011040) (cit. on p. 55).
- [75] Niklas Käming et al. “Unsupervised machine learning of topological phase transitions from experimental data”. In: *Machine Learning: Science and Technology* 2.3 (July 2021), p. 035037. DOI: [10.1088/2632-2153/abffe7](https://doi.org/10.1088/2632-2153/abffe7) (cit. on p. 55).
- [76] Hans-Jürgen Mikeska and Alexei K. Kolezhuk. “One-dimensional magnetism”. In: *Quantum Magnetism*. Ed. by Ulrich Schollwöck et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 1–83. DOI: [10.1007/BFb0119591](https://doi.org/10.1007/BFb0119591) (cit. on pp. 57, 58).
- [77] T. Giamarchi. *Quantum Physics in One Dimension*. Oxford, UK: Oxford University Press, 2004 (cit. on pp. 58, 63, 95, 115).
- [78] T. D. Kühner and H. Monien. “Phases of the one-dimensional Bose-Hubbard model”. In: *Phys. Rev. B* 58 (22 Dec. 1998), R14741–R14744. DOI: [10.1103/PhysRevB.58.R14741](https://doi.org/10.1103/PhysRevB.58.R14741) (cit. on pp. 58, 95).
- [79] Till D. Kühner, Steven R. White, and H. Monien. “One-dimensional Bose-Hubbard model with nearest-neighbor interaction”. In: *Phys. Rev. B* 61 (18 May 2000), pp. 12474–12489. DOI: [10.1103/PhysRevB.61.12474](https://doi.org/10.1103/PhysRevB.61.12474) (cit. on p. 58).
- [80] Maciej Lewenstein et al. “Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond”. In: *Advances in Physics* 56.2 (2007), pp. 243–379. DOI: [10.1080/00018730701223200](https://doi.org/10.1080/00018730701223200). eprint: <https://doi.org/10.1080/00018730701223200> (cit. on p. 58).
- [81] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. “Many-body physics with ultracold gases”. In: *Rev. Mod. Phys.* 80 (3 July 2008), pp. 885–964. DOI: [10.1103/RevModPhys.80.885](https://doi.org/10.1103/RevModPhys.80.885) (cit. on p. 58).
- [82] Erez Berg et al. “Rise and fall of hidden string order of lattice bosons”. In: *Phys. Rev. B* 77 (24 June 2008), p. 245119. DOI: [10.1103/PhysRevB.77.245119](https://doi.org/10.1103/PhysRevB.77.245119) (cit. on p. 58).
- [83] M. A. Cazalilla et al. “One dimensional bosons: From condensed matter systems to ultracold gases”. In: *Rev. Mod. Phys.* 83 (4 Dec. 2011), pp. 1405–1466. DOI: [10.1103/RevModPhys.83.1405](https://doi.org/10.1103/RevModPhys.83.1405) (cit. on pp. 58, 59).
- [84] Davide Rossini and Rosario Fazio. “Phase diagram of the extended Bose–Hubbard model”. In: *New Journal of Physics* 14.6 (June 2012), p. 065012. DOI: [10.1088/1367-2630/14/6/065012](https://doi.org/10.1088/1367-2630/14/6/065012) (cit. on pp. 58–60, 92, 94, 99).
- [85] Satoshi Ejima, Florian Lange, and Holger Fehske. “Spectral and Entanglement Properties of the Bosonic Haldane Insulator”. In: *Phys. Rev. Lett.* 113.2 (July 2014). DOI: [10.1103/physrevlett.113.020401](https://doi.org/10.1103/physrevlett.113.020401) (cit. on pp. 58–60).

- [86] Korbinian Kottmann et al. “Supersolid-superfluid phase separation in the extended Bose-Hubbard model”. In: *Phys. Rev. B* 104.17 (Nov. 2021). DOI: [10.1103/physrevb.104.174514](https://doi.org/10.1103/physrevb.104.174514) (cit. on pp. 58–60).
- [87] A. B. Kuklov and B. V. Svistunov. “Counterflow Superfluidity of Two-Species Ultracold Atoms in a Commensurate Optical Lattice”. In: *Phys. Rev. Lett.* 90 (10 Mar. 2003), p. 100401. DOI: [10.1103/PhysRevLett.90.100401](https://doi.org/10.1103/PhysRevLett.90.100401) (cit. on p. 61).
- [88] A. Kuklov, N. Prokof’ev, and B. Svistunov. “Commensurate Two-Component Bosons in an Optical Lattice: Ground State Phase Diagram”. In: *Phys. Rev. Lett.* 92 (5 Feb. 2004), p. 050402. DOI: [10.1103/PhysRevLett.92.050402](https://doi.org/10.1103/PhysRevLett.92.050402) (cit. on p. 61).
- [89] A. Hu et al. “Counterflow and paired superfluidity in one-dimensional Bose mixtures in optical lattices”. In: *Phys. Rev. A* 80 (2 Aug. 2009), p. 023619. DOI: [10.1103/PhysRevA.80.023619](https://doi.org/10.1103/PhysRevA.80.023619) (cit. on pp. 61, 63).
- [90] A. J. Leggett. “Topics in the theory of helium”. In: *Physica Fennica* 8 (1973), pp. 125–170 (cit. on p. 62).
- [91] A. Andreev and E. Bashkin. “Three-velocity hydrodynamics of superfluid solutions”. In: *Sov. Phys. JETP* 42.1 (1976), p. 164 (cit. on p. 62).
- [92] K. Sellin and E. Babaev. “Superfluid drag in the two-component Bose-Hubbard model”. In: *Phys. Rev. B* 97 (9 Mar. 2018), p. 094517. DOI: [10.1103/PhysRevB.97.094517](https://doi.org/10.1103/PhysRevB.97.094517) (cit. on p. 63).
- [93] Guansong Pang et al. “Deep Learning for Anomaly Detection”. In: *ACM Computing Surveys* 54.2 (Mar. 2021), pp. 1–38. DOI: [10.1145/3439950](https://doi.org/10.1145/3439950) (cit. on p. 65).
- [94] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach et al. Curran Associates, Inc., 2019, pp. 8024–8035 (cit. on pp. 69, 78).
- [95] The code is available together with the dataset at https://github.com/cerbero94/GAN_CP (cit. on p. 69).
- [96] Kaiming He et al. “Deep Residual Learning for Image Recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778. DOI: [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90) (cit. on p. 69).
- [97] Lian-Feng Dong et al. “Learning Deep Representations Using Convolutional Auto-Encoders with Symmetric Skip Connections”. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2018, pp. 3006–3010. DOI: [10.1109/ICASSP.2018.8462085](https://doi.org/10.1109/ICASSP.2018.8462085) (cit. on p. 69).
- [98] Andreas M. Läuchli. “Operator content of real-space entanglement spectra at conformal critical points”. 2013. arXiv: [1303.0741](https://arxiv.org/abs/1303.0741) [[cond-mat.stat-mech](https://arxiv.org/abs/1303.0741)] (cit. on pp. 69, 72, 74, 75, 91, 107).

- [99] Vincenzo Alba, Masudul Haque, and Andreas M. Läuchli. “Boundary-Locality and Perturbative Structure of Entanglement Spectra in Gapped Systems”. In: *Phys. Rev. Lett.* 108 (22 May 2012), p. 227201. DOI: [10.1103/PhysRevLett.108.227201](https://doi.org/10.1103/PhysRevLett.108.227201) (cit. on pp. 69, 91, 95, 97, 98).
- [100] Frank Pollmann et al. “Entanglement spectrum of a topological phase in one dimension”. In: *Phys. Rev. B* 81 (6 Feb. 2010), p. 064439. DOI: [10.1103/PhysRevB.81.064439](https://doi.org/10.1103/PhysRevB.81.064439) (cit. on pp. 70, 89).
- [101] Anushya Chandran et al. “Bulk-edge correspondence in entanglement spectra”. In: *Phys. Rev. B* 84 (20 Nov. 2011), p. 205136. DOI: [10.1103/PhysRevB.84.205136](https://doi.org/10.1103/PhysRevB.84.205136) (cit. on pp. 70, 89).
- [102] Xiao-Gang Wen. *Quantum Field Theory of Many-Body Systems: From the Origin of Sound to an Origin of Light and Electrons*. Oxford Scholarship Online, 2007. DOI: [DOI:10.1093/acprof:oso/9780199227259.001.0001](https://doi.org/10.1093/acprof:oso/9780199227259.001.0001) (cit. on pp. 70, 89).
- [103] Bei Zeng et al. *Quantum Information Meets Quantum Matter*. Springer New York, NY, 2019. DOI: [DOI:10.1007/978-1-4939-9084-9](https://doi.org/10.1007/978-1-4939-9084-9) (cit. on pp. 70, 89).
- [104] Xiaolong Deng and Luis Santos. “Entanglement spectrum of one-dimensional extended Bose-Hubbard models”. In: *Phys. Rev. B* 84 (8 Aug. 2011), p. 085138. DOI: [10.1103/PhysRevB.84.085138](https://doi.org/10.1103/PhysRevB.84.085138) (cit. on pp. 70, 89).
- [105] L. Lepori, G. De Chiara, and A. Sanpera. “Scaling of the entanglement spectrum near quantum phase transitions”. In: *Phys. Rev. B* 87 (23 June 2013), p. 235107. DOI: [10.1103/PhysRevB.87.235107](https://doi.org/10.1103/PhysRevB.87.235107) (cit. on pp. 70, 89).
- [106] Atsushi Ueda and Masaki Oshikawa. “Resolving the Berezinskii-Kosterlitz-Thouless transition in the two-dimensional XY model with tensor-network-based level spectroscopy”. In: *Phys. Rev. B* 104 (16 Oct. 2021), p. 165132. DOI: [10.1103/PhysRevB.104.165132](https://doi.org/10.1103/PhysRevB.104.165132) (cit. on p. 74).
- [107] M. Gerster et al. “Superfluid density and quasi-long-range order in the one-dimensional disordered Bose-Hubbard model”. In: *New Journal of Physics* 18 (2016). DOI: [10.1088/1367-2630/18/1/015015](https://doi.org/10.1088/1367-2630/18/1/015015) (cit. on pp. 75, 79, 95, 116).
- [108] Korbinian Kottmann et al. “Unsupervised mapping of phase diagrams of 2D systems from infinite projected entangled-pair states via deep anomaly detection”. In: *SciPost Phys.* 11 (2 2021), p. 25. DOI: [10.21468/SciPostPhys.11.2.025](https://doi.org/10.21468/SciPostPhys.11.2.025) (cit. on p. 79).
- [109] Paolo Zanardi, Marco Cozzini, and Paolo Giorda. “Ground state fidelity and quantum phase transitions in free Fermi systems”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2007.02 (Feb. 2007), p. L02002. DOI: [10.1088/1742-5468/2007/02/L02002](https://doi.org/10.1088/1742-5468/2007/02/L02002) (cit. on p. 82).

- [110] Ville Satopaa et al. “Finding a ”Kneedle” in a Haystack: Detecting Knee Points in System Behavior”. In: *2011 31st International Conference on Distributed Computing Systems Workshops*. 2011, pp. 166–171. DOI: [10.1109/ICDCSW.2011.20](https://doi.org/10.1109/ICDCSW.2011.20) (cit. on pp. 85, 86).
- [111] H. Francis Song, Stephan Rachel, and Karyn Le Hur. “General relation between entanglement and fluctuations in one dimension”. In: *Phys. Rev. B* 82 (1 July 2010), p. 012405. DOI: [10.1103/PhysRevB.82.012405](https://doi.org/10.1103/PhysRevB.82.012405) (cit. on pp. 89, 90, 99).
- [112] Brian Skinner, Jonathan Ruhman, and Adam Nahum. “Measurement-Induced Phase Transitions in the Dynamics of Entanglement”. In: *Phys. Rev. X* 9 (3 July 2019), p. 031009. DOI: [10.1103/PhysRevX.9.031009](https://doi.org/10.1103/PhysRevX.9.031009) (cit. on p. 89).
- [113] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. “Measurement-driven entanglement transition in hybrid quantum circuits”. In: *Phys. Rev. B* 100 (13 Oct. 2019), p. 134306. DOI: [10.1103/PhysRevB.100.134306](https://doi.org/10.1103/PhysRevB.100.134306) (cit. on p. 89).
- [114] A. J. Daley et al. “Measuring Entanglement Growth in Quench Dynamics of Bosons in an Optical Lattice”. In: *Phys. Rev. Lett.* 109 (2 July 2012), p. 020505. DOI: [10.1103/PhysRevLett.109.020505](https://doi.org/10.1103/PhysRevLett.109.020505) (cit. on p. 90).
- [115] Rajibul Islam et al. “Measuring entanglement entropy in a quantum many-body system”. In: *Nature* 528.7580 (Dec. 2015), pp. 77–83. DOI: [10.1038/nature15750](https://doi.org/10.1038/nature15750) (cit. on p. 90).
- [116] Adam M. Kaufman et al. “Quantum thermalization through entanglement in an isolated many-body system”. In: *Science* 353.6301 (2016), pp. 794–800. DOI: [10.1126/science.aaf6725](https://doi.org/10.1126/science.aaf6725) (cit. on p. 90).
- [117] Tiff Brydges et al. “Probing Rènyi entanglement entropy via randomized measurements”. In: *Science* 364.6437 (2019), pp. 260–263. DOI: [10.1126/science.aau4963](https://doi.org/10.1126/science.aau4963) (cit. on p. 90).
- [118] G. Giudici et al. “Entanglement Hamiltonians of lattice models via the Bisognano-Wichmann theorem”. In: *Phys. Rev. B* 98 (13 Oct. 2018), p. 134403. DOI: [10.1103/PhysRevB.98.134403](https://doi.org/10.1103/PhysRevB.98.134403) (cit. on p. 90).
- [119] Zhanyu Ma et al. “Symmetric inseparability and number entanglement in charge-conserving mixed states”. In: *Phys. Rev. A* 105 (4 Apr. 2022), p. 042416. DOI: [10.1103/PhysRevA.105.042416](https://doi.org/10.1103/PhysRevA.105.042416) (cit. on p. 90).
- [120] Alexander Lukin et al. “Probing entanglement in a many-body-localized system”. In: *Science* 364.6437 (2019), pp. 256–260. DOI: [10.1126/science.aau0818](https://doi.org/10.1126/science.aau0818) (cit. on pp. 90, 91).
- [121] Maximilian Kiefer-Emmanouilidis et al. “Evidence for Unbounded Growth of the Number Entropy in Many-Body Localized Phases”. In: *Phys. Rev. Lett.* 124 (24 June 2020), p. 243601. DOI: [10.1103/PhysRevLett.124.243601](https://doi.org/10.1103/PhysRevLett.124.243601) (cit. on p. 90).

- [122] Maximilian Kiefer-Emmanouilidis et al. “Unlimited growth of particle fluctuations in many-body localized phases”. In: *Annals of Physics* 435 (2021). Special Issue on Localisation 2020, p. 168481. DOI: <https://doi.org/10.1016/j.aop.2021.168481> (cit. on p. 90).
- [123] H. Francis Song et al. “Bipartite fluctuations as a probe of many-body entanglement”. In: *Phys. Rev. B* 85 (3 Jan. 2012), p. 035409. DOI: [10.1103/PhysRevB.85.035409](https://doi.org/10.1103/PhysRevB.85.035409) (cit. on pp. 90, 91, 95, 115, 117).
- [124] Stephan Rachel et al. “Detecting Quantum Critical Points Using Bipartite Fluctuations”. In: *Phys. Rev. Lett.* 108 (11 Mar. 2012), p. 116401. DOI: [10.1103/PhysRevLett.108.116401](https://doi.org/10.1103/PhysRevLett.108.116401) (cit. on pp. 90, 116).
- [125] Israel Klich and Leonid Levitov. “Quantum Noise as an Entanglement Meter”. In: *Phys. Rev. Lett.* 102 (10 Mar. 2009), p. 100502. DOI: [10.1103/PhysRevLett.102.100502](https://doi.org/10.1103/PhysRevLett.102.100502) (cit. on p. 90).
- [126] Moshe Goldstein and Eran Sela. “Symmetry-Resolved Entanglement in Many-Body Systems”. In: *Phys. Rev. Lett.* 120 (20 May 2018), p. 200602. DOI: [10.1103/PhysRevLett.120.200602](https://doi.org/10.1103/PhysRevLett.120.200602) (cit. on pp. 90, 91).
- [127] Riccarda Bonsignori, Paola Ruggiero, and Pasquale Calabrese. “Symmetry resolved entanglement in free fermionic systems”. In: *Journal of Physics A: Mathematical and Theoretical* 52.47 (Oct. 2019), p. 475302. DOI: [10.1088/1751-8121/ab4b77](https://doi.org/10.1088/1751-8121/ab4b77) (cit. on p. 90).
- [128] Luca Capizzi, Paola Ruggiero, and Pasquale Calabrese. “Symmetry resolved entanglement entropy of excited states in a CFT”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2020.7 (July 2020), p. 073101. DOI: [10.1088/1742-5468/ab96b6](https://doi.org/10.1088/1742-5468/ab96b6) (cit. on p. 90).
- [129] Pasquale Calabrese et al. “Full counting statistics in the gapped XXZ spin chain”. In: *EPL (Europhysics Letters)* 129.6 (Apr. 2020), p. 60007. DOI: [10.1209/0295-5075/129/60007](https://doi.org/10.1209/0295-5075/129/60007) (cit. on p. 90).
- [130] Nicolas Laflorencie and Stephan Rachel. “Spin-resolved entanglement spectroscopy of critical spin chains and Luttinger liquids”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2014.11 (Nov. 2014), P11013. DOI: [10.1088/1742-5468/2014/11/p11013](https://doi.org/10.1088/1742-5468/2014/11/p11013) (cit. on p. 91).
- [131] Korbinian Kottmann et al. “Supersolid-superfluid phase separation in the extended Bose-Hubbard model”. In: *Phys. Rev. B* 104 (17 Nov. 2021), p. 174514. DOI: [10.1103/PhysRevB.104.174514](https://doi.org/10.1103/PhysRevB.104.174514) (cit. on p. 94).
- [132] Wei Chen, Kazuo Hida, and B. C. Sanctuary. “Ground-state phase diagram of $S = 1$ XXZ chains with uniaxial single-ion-type anisotropy”. In: *Phys. Rev. B* 67 (10 Mar. 2003), p. 104401. DOI: [10.1103/PhysRevB.67.104401](https://doi.org/10.1103/PhysRevB.67.104401) (cit. on p. 98).
- [133] Erez Berg et al. “Rise and fall of hidden string order of lattice bosons”. In: *Phys. Rev. B* 77 (24 June 2008), p. 245119. DOI: [10.1103/PhysRevB.77.245119](https://doi.org/10.1103/PhysRevB.77.245119) (cit. on p. 98).

- [134] Satoshi Ejima, Florian Lange, and Holger Fehske. “Spectral and Entanglement Properties of the Bosonic Haldane Insulator”. In: *Phys. Rev. Lett.* 113 (2 July 2014), p. 020401. DOI: [10.1103/PhysRevLett.113.020401](https://doi.org/10.1103/PhysRevLett.113.020401) (cit. on p. 98).
- [135] Abubakar Abid, Muhammad Fatih Balin, and James Zou. *Concrete Autoencoders for Differentiable Feature Selection and Reconstruction*. 2019. DOI: [10.48550/ARXIV.1901.09346](https://doi.org/10.48550/ARXIV.1901.09346). eprint: [1901.09346](https://arxiv.org/abs/1901.09346) (cit. on pp. 101, 102).
- [136] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. *The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables*. 2016. DOI: [10.48550/ARXIV.1611.00712](https://doi.org/10.48550/ARXIV.1611.00712). eprint: [1611.00712](https://arxiv.org/abs/1611.00712) (cit. on p. 103).

Acknowledgements

I would like to express my gratitude to my supervisors Ale, Teo and Elisa for their guidance throughout this project. With them, I always felt free to express my feelings, doubts or enthusiasm in tackling the research work and I could always rely on their support. I would also like to thank Enrico and Andreas for their help and the very interesting discussions. Many thanks also go to the research team at the FZJ – and again to Teo for making this possible – for welcoming and hosting me during my time in Germany and making it a wonderful experience both personally and academically. I am truly thankful to Luciano who believed in me and made me part of the new adventure with Travelbrain.

A big thank goes to my colleagues in Trento: Piero, Piero, Alberto, Sergio, Samu and Lorenzo. I am grateful for all the time we spent together, the discussions, the pizzas in Mensa, the coffees at the vending machine, the skiing, the nights out. I thank my colleagues in Germany: Francesco, Daniel, Nicklas, Michael, Markus and Martino.

A special thank needs to go to my friends Yle, Fede, Enzo with whom I have shared so much since I started my journey in Trento, for their friendship and open doors. I would like to extend my grateful thanks to all the Aachen Gang: Dani, Moritz, Nihar, Nicoletta, Guille, Paco, Alice, Erik, Natalia and Jordi who have made Aachen home. I thank Liuc, AleF and my brothers Nello and Gio for their support and for always being there despite the distance. I truly appreciate all of you.

I thank Dad and Mom because they have always been my rocks and my reference point, Giorgia and Lorenzo, my nana Anna, who is the strongest person I know, and grandpas Angelo and Marietta who have always lovingly taken care of me.

The biggest thank-you goes to my love AleFaghe who has been the best I could ever had on my side for this long and wonderful journey. She simply made everything possible.