

Article

Variational Autoencoder Reconstruction of Complex Many-Body Physics

Ilia A. Luchnikov^{1,2}, Alexander Ryzhov¹ , Pieter-Jan Stas³ and Sergey N. Filippov^{2,4,5}  and Henni Ouerdane^{1,*} 

¹ Center for Energy Science and Technology, Skolkovo Institute of Science and Technology, 3 Nobel Street, Skolkovo, Moscow Region 121205, Russia; Ilia.Luchnikov@skoltech.ru (I.A.L.); a.ryzhov@skoltech.ru (A.R.)

² Moscow Institute of Physics and Technology, Institutskii Per. 9, Dolgoprudny, Moscow Region 141700, Russia; sergey.filippov@phystech.edu

³ Department of Applied Physics, Stanford University 348 Via Pueblo Mall, Stanford, CA 94305, USA; pjstas@stanford.edu

⁴ Valiev Institute of Physics and Technology of Russian Academy of Sciences, Nakhimovskii Pr. 34, Moscow 117218, Russia

⁵ Steklov Mathematical Institute of Russian Academy of Sciences, Gubkina St. 8, Moscow 119991, Russia

* Correspondence: h.ouerdane@skoltech.ru

Received: 9 October 2019; Accepted: 6 November 2019; Published: 7 November 2019



Abstract: Thermodynamics is a theory of principles that permits a basic description of the macroscopic properties of a rich variety of complex systems from traditional ones, such as crystalline solids, gases, liquids, and thermal machines, to more intricate systems such as living organisms and black holes to name a few. Physical quantities of interest, or equilibrium state variables, are linked together in equations of state to give information on the studied system, including phase transitions, as energy in the forms of work and heat, and/or matter are exchanged with its environment, thus generating entropy. A more accurate description requires different frameworks, namely, statistical mechanics and quantum physics to explore in depth the microscopic properties of physical systems and relate them to their macroscopic properties. These frameworks also allow to go beyond equilibrium situations. Given the notably increasing complexity of mathematical models to study realistic systems, and their coupling to their environment that constrains their dynamics, both analytical approaches and numerical methods that build on these models show limitations in scope or applicability. On the other hand, machine learning, i.e., data-driven, methods prove to be increasingly efficient for the study of complex quantum systems. Deep neural networks, in particular, have been successfully applied to many-body quantum dynamics simulations and to quantum matter phase characterization. In the present work, we show how to use a variational autoencoder (VAE)—a state-of-the-art tool in the field of deep learning for the simulation of probability distributions of complex systems. More precisely, we transform a quantum mechanical problem of many-body state reconstruction into a statistical problem, suitable for VAE, by using informationally complete positive operator-valued measure. We show, with the paradigmatic quantum Ising model in a transverse magnetic field, that the ground-state physics, such as, e.g., magnetization and other mean values of observables, of a whole class of quantum many-body systems can be reconstructed by using VAE learning of tomographic data for different parameters of the Hamiltonian, and even if the system undergoes a quantum phase transition. We also discuss challenges related to our approach as entropy calculations pose particular difficulties.

Keywords: complex systems thermodynamics; machine learning; quantum phase transition; Ising model; variational autoencoder

1. Introduction

The development of the dynamical theory of heat or classical equilibrium thermodynamics as we know it was possible only with empirical data collection, processing, and analysis, which led, through a phenomenological approach, to the definition of two fundamental physical concepts, the actual pillars of the theory: energy and entropy [1]. It is with these two concepts that the laws (or principles) of thermodynamics could be stated and the absolute temperature be given a first proper definition. Though energy remains as fully enigmatic as entropy from the ontological viewpoint, the latter concept is not completely understood from the physical viewpoint. This of course did not preclude the success of equilibrium thermodynamics as evidenced not only by the development of thermal sciences and engineering, but also because of its cognate fields that owe it, at least partly or as an indirect consequence, their birth, from quantum physics to information theory.

Early attempts to refine and give thermodynamics solid grounds started with the development of the kinetic theory of gases and of statistical physics, which in turn permitted studies of irreversible processes with the development of nonequilibrium thermodynamics [2–6] and later on finite-time thermodynamics [7–9], thus establishing closer ties between the concrete notion of irreversibility and the more abstract entropy, notably with Boltzmann's statistical definition [10] and Gibbs' ensemble theory [11]. Notwithstanding conceptual difficulties inherent to the foundations of statistical physics, such as, e.g., irreversibility and the ergodic hypothesis [12,13], entropy acquired a meaningful statistical character and the scope of its definitions could be extended beyond thermodynamics, thus paving the way to information theory, as information content became a physical quantity per se, i.e., something that can be measured [14]. Additionally, although quantum physics developed independently from thermodynamics, it extended the scope of statistical physics with the introduction of quantum statistics, led to the definition of the von Neumann entropy [15], and also introduced new problems related to small, i.e., mesoscopic and nanoscopic systems [16,17], down to nuclear matter [18], where the concepts of thermodynamic limit and ensuing standard definitions of thermodynamic quantities may be put at odds.

Quantum physics problems that overlap with thermodynamics are typically classified into different categories: ground state characterization [19], thermal state characterization at finite temperature [20], the so-called eigenstate thermalization hypothesis [21–25], calculation of the dynamics of either closed or open systems [26,27], state reconstruction from tomographic data [28], and quantum system control, which, given the complexity for its implementation, requires the development of new methods [29]. Among the rich variety of methods applicable to such problems, including, e.g., mean-field approach [30], slave particle approach [31], dynamical mean-field theory [32], nonperturbative methods based on functional integrals [33], we believe two large families of techniques are of particular interest for numerical studies of many-body systems when strong correlations must be accounted for: One is based on the quantum Monte Carlo (QMC) framework [34], which is powerful to overcome the curse of dimensionality by using the stochastic estimation of high-dimensional integrals; the other family encompasses methods that search solutions in the parametric set of functions, also called ansatz. The most used ansatzes are based on different tensor network architectures [35,36] as tensor network-based methods show state-of-the-art performance for the characterization of one-dimensional strongly correlated quantum systems. One can solve either the ground-state problem by using the variational matrix product state (MPS) ground state search [37] or a dynamical problem using a time-evolving block decimation (TEBD) algorithm [38]. Quantum criticality of one-dimensional systems also can be studied by using a more advanced architecture called multiscale entanglement renormalization ansatz (MERA) [39]. The application of tensor networks is not restricted to one-dimensional systems, and one can describe an open quantum dynamics [40], characterize the numerical complexity of an open quantum dynamics [41,42], perform tomography of non-Markovian quantum processes by using tensor networks [43,44], analyze properties of two dimensional quantum lattices by using projected entangled pair states (PEPS) [45], or solve classical statistical physics problems [46,47].

The cross-fertilization of quantum physics and thermodynamics has benefited much from the powerful quantum formalism and computational techniques; however, as thermodynamic concepts evolved from intuitive/phenomenological definitions to classical-mechanics constructs, extended with quantum physics and formalism when needed, thermodynamics, in spite of its undeniable theoretical and practical successes, never managed to fully mature into a genuine fundamental theory that firmly rests on strong basic postulates. On one hand, this led a growing number of physicists to consider thermodynamics as incomplete, and on the other, to think quantum theory as the underlying framework from which equilibrium and nonequilibrium thermodynamics emerge. Quantum thermodynamics [48,49] is a fairly recent field of play, where new ideas are tested while revisiting old problems related to cycles, engines, refrigerators, and entropy production, to name a few [50,51]. Further, quantum technology is a burgeoning field at the interface of physics and engineering, which seeks to develop devices able to harness quantum effects for computing and secure communication purposes [52,53]. The wide scale development of such a kind of systems, which irreversibly interact with an infinite environment, rests on the ability to properly simulate the open quantum dynamics of their many-body properties and analyze coherence and dissipation at the quantum level.

How fast quantum thermodynamics will progress is difficult to anticipate as there exist numerous unsolved problems, especially those related to the proper characterization of the physical processes, e.g., what qualifies as heat or work on ultrashort time and length scales, where averages become irrelevant is unclear, and how the laws of thermodynamics may be systematically adapted still may be debated. To mitigate risks of slow progress, one may resort to approaches that do not rely on models of systems, but rather on data, the idea being to gain actual knowledge and understanding from data irrespective of how complex the studied system is. Machine learning (ML) provides perfectly suited tools for that purpose [54]. ML has a rather long history that can be dated back with the works of Bayes (1763) on prior knowledge that can be used to calculate the probability of an event as formulated by Laplace (1812). Much later (1913), Markov chains were proposed as a tool to describe sequences of events, each being characterized by a probability of occurrence that depends on the actuality of the previous event only. The main milestone is in 1950, with Turing's machine that can learn [55], shortly followed in 1951 by the first neural network machine [56]. Thanks to the huge increase in computational power over the last two decades, ML is now used for a wide variety of problems [54], and quantum machine learning now shows extraordinary potential for faster and more efficient than ever treatment of complex quantum systems problems [57], one major challenge still residing in the development of the hardware capable to harness and transform this potentiality into actual tool.

With the recent success in the field of deep learning, tools other than those based on tensor networks work as well as an ansatz. Restricted Boltzmann machine has been successfully applied as an ansatz to a ground state search, dynamics calculation, and quantum tomography [58–60], as well as convolution neural network to the two-dimensional frustrated $J_1 - J_2$ model [61]. The deep autoregressive model was applied very efficiently and elegantly to a ground state search of many-body quantum system and to classical statistical physics as well [62,63]. It was also recently shown how ML can establish and classify with high accuracy the chaotic or regular behavior of quantum billiards models and XXZ spin chains [64]. Thus, it can be useful to transfer deep architectures from the field of deep learning to the area of many-body quantum systems. A variational autoencoder (VAE) was used for sampling from probability distributions of quantum states in [65]; in the present work, we show that state-of-the-art generative architecture called conditional VAE can be applied to describe the whole family of the ground states of a quantum many-body system. For that purpose, using quantum tomography (albeit in an approximate fashion as discussed below) and reconstruction tools developed in [66], we consider the paradigmatic Ising model in a transverse-field as an illustration of the usefulness and efficiency of our approach. The use of VAE in such a problem is justified by the simplicity of VAE training, as well as its expressibility [67].

The article is organized as follows. In Section 2, we give a brief recap of the physics of the Ising model in a transverse field. In Section 3, we develop our generative model in the framework of the tensor network. Section 4 is devoted to the variational autoencoder architecture. The results are shown and discussed in Section 5. The article ends with concluding remarks, followed by a short series of appendices.

2. Transverse-Field Ising Model

Among the rich variety of condensed matter systems, magnetic materials are a source of many fruitful problems, whose studies and solutions inspired discussions and new models beyond their immediate scope. The Kondo effect (existence of a minimum of electrical resistivity at low temperature in metals due to the presence of magnetic impurities) is one such problem [68,69], as it provides an excellent basis for studies of quantum criticality and absolute zero-temperature phase transitions [70,71] and, also, on a more fundamental level, a concrete example of asymptotic freedom [69]. Assuming infinite on-site repulsion, the single-impurity Anderson model [68,72] was used to establish a correspondence between Hamiltonian language and path integral for the development of nonperturbative methods in quantum field theory [73,74]. One other important model is that of the Heisenberg Hamiltonian, defined for the study of ferromagnetic materials, and which, assuming a crystal subjected to an external magnetic field \mathbf{B} , reads [75] as

$$H = - \sum_{\langle i,j \rangle} J_{ij} \hat{S}^i \hat{S}^j - \mathbf{h} \cdot \sum_j \hat{S}^j \quad (1)$$

where, for ease of notations, we introduced $\mathbf{h} = g\mu_B\mathbf{B}$, with g being the Landé factor and $\mu_B = e\hbar/2m_e$ being the Bohr magneton (e : elementary electric charge, and m_e : electron mass); J_{ij} is a parameter that characterizes the nearest-neighbors exchange interaction between electron spins on the crystal sites i and j (the quantum spins \hat{S}^i and \hat{S}^j are vector operators whose components are proportional to the Pauli matrices). For simplicity, one may consider $J_{ij} \equiv J$ constant. If $J > 0$, then the system is ferromagnetic and if $J < 0$ the system is antiferromagnetic. Hereafter, we fix the electron's magnetic moment $g\mu_B = 1$.

Although Equation (1) has a fairly simple form, the exact calculation of the partition function is

$$Z = \text{Tr} e^{-\beta H} \quad (2)$$

where $\beta = 1/k_B T$ is the inverse thermal energy, which is possible on the analytical level with the mean-field approximation that simplifies the Hamiltonian (1), and also for one-dimensional systems, one difficulty of the Heisenberg Hamiltonian being that the three components of a spin vector operator do not commute. That said, Heisenberg's Hamiltonian is very useful to, e.g., study spin frustration [76], entanglement entropy [77], and also serve as a test case for density-matrix renormalization group algorithms [78]. Under zero field, Heisenberg's Hamiltonian is also a simplified form of the Hubbard model at half-filling, thus including ferromagnetism in the scope of strongly correlated systems studies.

A particular, but very important, approximation of Heisenberg's Hamiltonian, whose significance lies in physics, especially for the study of critical phenomena, cannot be underestimated: the so-called Ising model. In its initial formulation [79], Ising spins are N classical variables, which may take ± 1 as values and form a one-dimensional (1D) system characterized by free or periodic boundary conditions. The classical partition function Z may be calculated analytically for the 1D Ising model, and quantities such as the average total magnetization are obtained directly [80]:

$$M = \frac{1}{\beta} \frac{\partial \ln Z}{\partial h} \quad (3)$$

In the present work, we consider a 1D quantum spin chain whose Hilbert space is given by $\mathcal{H} = \otimes_i^N \mathbb{C}^2$. The system is described by the transverse-field Ising (TFI) Hamiltonian [81]:

$$H = -J \sum_{\langle i,j \rangle} \sigma_z^i \sigma_z^{i+1} + h_x \sum_{i=1}^N \sigma_x^i. \quad (4)$$

where σ_α^i ($\alpha \equiv x, z$) is the Pauli matrix for the α -component of the i -th spin in the chain, and h_x is the magnetic field applied in the transverse direction x . In this case, the spins are no longer the classical Ising ones and the two terms that compose the Hamiltonian H do not commute, therefore requiring a full quantum approach. An example of a real-world system that may be studied as a quantum Ising chain is cobalt niobate (CoNb_2O_6); in this case, the spins that undergo the phase transition as the transverse field varies are those of the Co^{2+} ions [82]. The spin states are denoted $|+\rangle_i$ and $|-\rangle_i$ at ion site i . There are two possible ground states: when all N spins are in the state $|+\rangle$ or in the state $|-\rangle$, i.e., when they are all aligned, which defines the ferromagnetic phase.

The phase transition from the ferromagnetic phase to the paramagnetic phase that we speak of now is of a quantum nature, and not of a thermal nature, as here it is driven only by the external magnetic field. More precisely, when the transverse field h_x is applied with sufficient strength, the spins align along the x direction, and the spin state at site i is given as the superposition $(|+\rangle_i + |-\rangle_i) / \sqrt{2}$, which is nothing else but the eigenstate of the x -component of the spin. Therefore, in this particular case, there is no need to raise the temperature of the system initially in the ferromagnetic phase beyond the Curie temperature to make it a paramagnet: the many-body system remains in its ground state, but its properties have changed. Further, note that unlike for the ferromagnetic phase, the quantum paramagnetic phase has spin-inversion symmetry. An insightful discussion on quantum criticality can be found in Reference [83].

Now, we briefly comment on the quantity $\beta = 1/k_B T$ in the context of quantum phase transitions, which, strictly speaking, can only occur at temperature $T = 0$ K. In fact, close to the absolute zero, where $\beta \rightarrow \infty$, their signatures can be observed as quantum fluctuations dominate thermal fluctuations in the criticality region, where the quantum critical point lies. The imaginary time formalism [84], where $\exp(-\beta H)$ is interpreted as an evolution operator, and the partition function Z as a path integral, provides a way to map a quantum problem onto a classical one with the introduction of the imaginary time β resulting from a Wick rotation in the complex plane, thus yielding one extra dimension to the model. In classical thermodynamics, to observe a phase transition in a system requires that its size (i.e., the number of constituents N) tends to infinity so that the order parameter is non-analytic at the transition point; so, for the quantum transition, the thermodynamic limit entails the limit $\beta \rightarrow \infty$ also: the 1D TFI model is mapped onto an equivalent 2D classical Ising model [85]. The imaginary time formalism permits implementation of classical Monte Carlo simulations to study quantum systems. Further discussion, including the sign problem for the quantum spin-1/2 system, is available in Reference [4].

We have chosen the transverse-field Ising model as an illustrative case for our study for several reasons. First, as this system is 1-dimensional, we can apply an MPS variational ground state solver [37], and therefore obtain the ground state solution in MPS representation. We can then perform fast and exact sampling for generation of large data sets for the training of the VAE. Next, this model can be solved analytically, which allows us to adequately benchmark our results. Finally, this model shows a nontrivial behavior around the quantum phase transition point at $h_x = 1$, and thus constitutes an interesting example to apply a VAE.

3. Generative Model as a Quantum State

Many-body quantum physics is rich in high-dimensional problems. Often, however, with increasing dimensionality, these become extremely difficult or impossible to solve. One solving

method is through the reformulation of the quantum mechanical problem as a statistical problem, when possible. This way, machine learning can be used to effectively solve such a problem, as machine learning is a tool for the solving of high-dimensional statistical problems [86]. Probabilistic interpretation allows for using powerful sampling-based methods that work efficiently with high dimensional data.

An example of the reformulation of a quantum problem as a statistical problem is with informationally complete (IC) positive-operator valued measures (POVMs) [87]. POVMs describe the most general measurements of a quantum system. Each particular POVM is defined by a set of positive semidefinite operators M^α , with the normalization condition $\sum_\alpha M^\alpha = \mathbb{1}$, where $\mathbb{1}$ is the identity operator. The fact that the POVM is informationally complete means that using measurement outcomes one can reconstruct the state of a system with arbitrary accuracy.

The probability of measurement outcome for a quantum system with the density operator ρ is governed by Born’s rule: $P[\alpha] = \text{Tr}(\rho M^\alpha)$, where $\{M^\alpha\}$ is a particular POVM and α is an outcome result. In other words, any density matrix can be mapped on a mass function, although not all mass functions can be mapped on a density matrix [88,89]. Some mass functions lead to non-positive semidefinite “density matrices”, which is not physically allowed. As such, quantum theory is a constrained version of probability theory. For a many-body system, these constraints can be very complicated, and direct consideration of quantum theory as a constrained probability theory is not fruitful. However, if one can access the samples of the IC POVM induced mass function, which is by definition physically allowed, this mass function can be reconstructed using generative modeling [66,67]. Samples can be obtained either by performing generalized measurements over the quantum system or by in silico simulation.

In the present work, we simulate measurements of the ground state of a spin chain with the TFI Hamiltonian, Equation (4). As a local (one spin) IC POVM, we use the so-called symmetric IC POVM for qubits (tetrahedral) POVM [90]:

$$M_{\text{tetra}}^\alpha = \frac{1}{4} (\mathbb{1} + s^\alpha \sigma), \quad \alpha \in (0, 1, 2, 3), \quad \sigma = (\sigma_x, \sigma_y, \sigma_z),$$

$$s^0 = (0, 0, 1), \quad s^1 = \left(\frac{2\sqrt{2}}{3}, 0, -\frac{1}{3}\right), \quad s^2 = \left(-\frac{\sqrt{2}}{3}, \sqrt{\frac{2}{3}}, -\frac{1}{3}\right), \quad s^3 = \left(-\frac{\sqrt{2}}{3}, -\sqrt{\frac{2}{3}}, -\frac{1}{3}\right). \quad (5)$$

Note that the many-spin generalization of local IC POVM can easily be obtained by considering the tensor product of local ones:

$$M_{\text{tetra}}^{\alpha_1, \dots, \alpha_N} = M_{\text{tetra}}^{\alpha_1} \otimes M_{\text{tetra}}^{\alpha_2} \otimes \dots \otimes M_{\text{tetra}}^{\alpha_N}. \quad (6)$$

To simulate measurements outcome under the IC POVM described above, we implement the following numerical scheme: First, we run a variational MPS ground state solver to obtain the ground state of the TFI model in the MPS form:

$$\Omega_{i_1, i_2, \dots, i_N} = \sum_{\beta_1, \beta_2, \dots, \beta_{N-1}} A_{i_1 \beta_1}^1 A_{\beta_1 i_2 \beta_2}^2 \dots A_{\beta_{N-1} i_N}^N \quad (7)$$

where we use the tensor notation instead of the bra-ket notation for further simplicity, and we obtain the MPS representation of IC POVM induced mass function:

$$P[\alpha_1, \alpha_2, \dots, \alpha_N] = \sum_{\delta_1, \delta_2, \dots, \delta_{N-1}} \pi_{\alpha_1 \delta_1} \pi_{\delta_1 \alpha_2 \delta_2} \dots \pi_{\delta_{N-1} \alpha_N}$$

$$\pi_{\delta_{n-1} \alpha_n \delta_n} = \pi \underbrace{\beta_{n-1} \beta'_{n-1}}_{\text{multi-index } \delta_{n-1}} \alpha_n \underbrace{\beta_n \beta'_n}_{\text{multi-index } \delta_n} = [M_{\text{tetra}}]_{ij}^{\alpha_n} A_{\beta_{n-1} j \beta_n}^n [A^n]_{\beta'_{n-1} i \beta'_n}^* \quad (8)$$

whose diagrammatic representation [35] is shown in Figure 1. Next, we produce a set of samples of size M : $\{\alpha_1^i, \alpha_2^i, \dots, \alpha_N^i\}_{i=1}^M$ from the given probability. The sampling can be efficiently implemented

as shown in Appendix B. We call this set of samples (outcome measurements) a data set, which may then be used to train a generative model $p[\alpha_1, \alpha_2, \dots, \alpha_N | \theta]$ to emulate the true mass function $P[\alpha_1, \alpha_2, \dots, \alpha_N]$. Here, θ is the set of parameters of the generative model, which is trained by maximizing the logarithmic likelihood $\mathcal{L}(\theta) = \sum_{i=1}^M \log p[\alpha_1^i, \alpha_2^i, \dots, \alpha_N^i | \theta]$ with respect to the parameters θ [91]. The trained generative model fully characterizes a quantum state. The density matrix is obtained by applying an inverse transformation to the mass function [92]:

$$\begin{aligned} \rho &= \sum_{\alpha_1, \alpha_2, \dots, \alpha_N} p[\alpha_1, \alpha_2, \dots, \alpha_N | \theta] [M_{\text{tetra}}^{\alpha_1}]^{-1} \otimes [M_{\text{tetra}}^{\alpha_2}]^{-1} \otimes \dots \otimes [M_{\text{tetra}}^{\alpha_N}]^{-1}, \\ [M_{\text{tetra}}^{\alpha}]^{-1} &= \sum_{\alpha'} T_{\alpha\alpha'}^{-1} M_{\text{tetra}}^{\alpha'}, \\ T_{\alpha\alpha'} &= \text{Tr} \left(M_{\text{tetra}}^{\alpha} M_{\text{tetra}}^{\alpha'} \right), \end{aligned} \tag{9}$$

the diagrammatic representation of which is given in Figure 2. Note that the summation included in the density matrix representation is numerically intractable, but we can estimate it using samplings from the generative model.

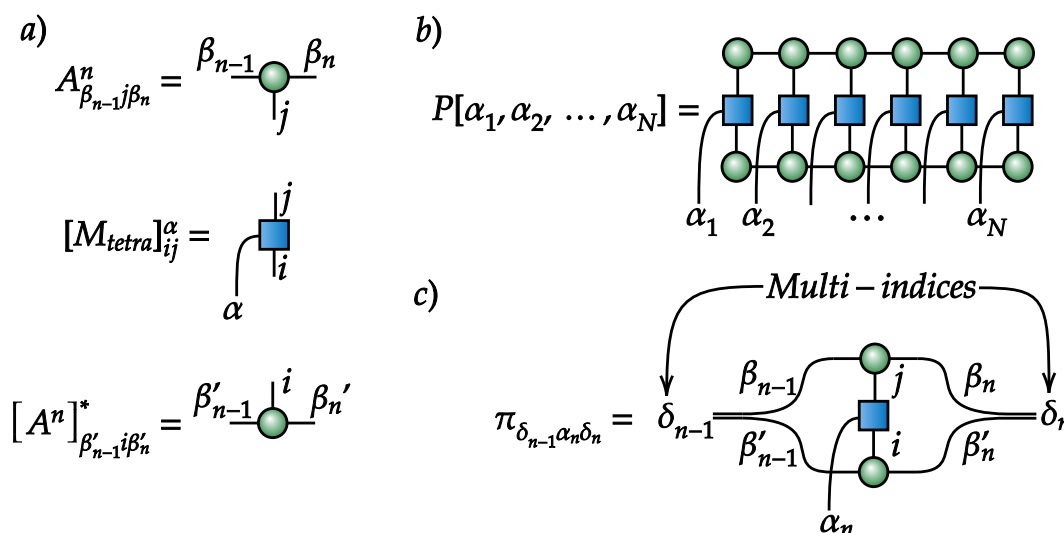


Figure 1. Tensor diagrams for (a) building blocks, (b) matrix product state (MPS) representation of measurement outcome probability, and (c) its subtensor.

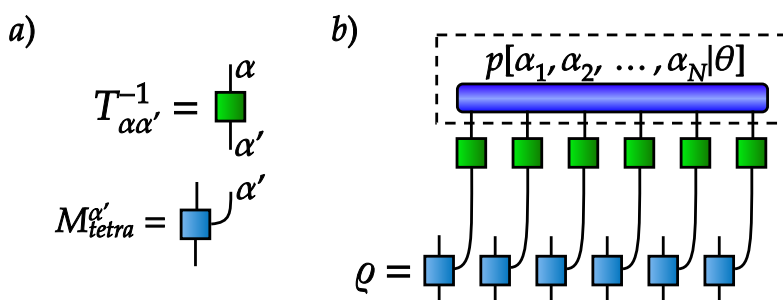


Figure 2. Tensor diagrams for (a) building blocks and (b) inverse transformation from a mass function to a density matrix.

Our goal is to use a generative model as an effective representation of quantum states to calculate the mean values of observables such as, e.g., two-point and higher-order correlation functions. An explicit expression of the two-point correlation function obtained by sampling from the trained generative model is shown in Figure 3. To obtain the ground state of the TFI model, we use a

variational MPS ground state search, and we pick the bond dimension of MPS equal to 25 and perform 5 DMRG sweeps to get an approximate ground state in the MPS form. We use the variational MPS solver provided by the mpnum toolbox [93].

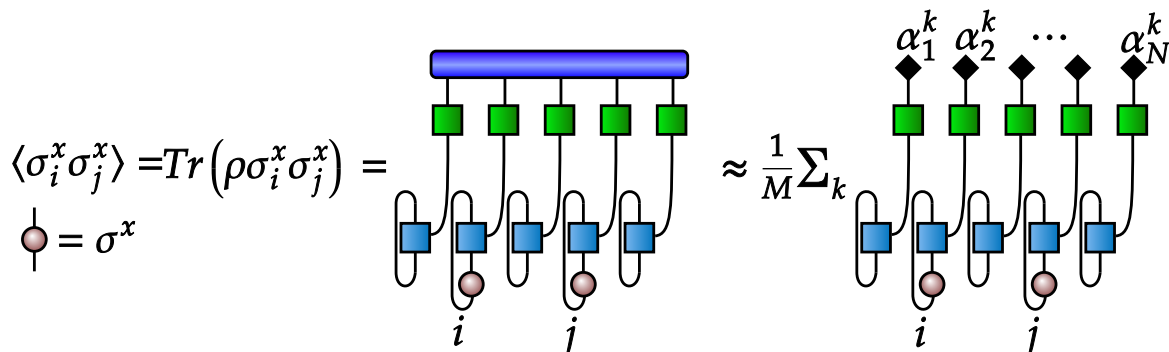


Figure 3. Tensor diagrams representing calculation of two-point correlation function.

4. Variational Autoencoder Architecture

In our work, we use a conditional VAE [94] to represent quantum states. A conditional VAE is a generative model expressed by the following probability distribution,

$$p[x|\theta, h] = \int p[x|z, \theta, h]p[z]dz, \tag{10}$$

where x is the data we want to simulate; θ represents the VAE parameters, which can be tuned to get the desired probability distribution over x ; h is the condition; and z is a vector of latent variables. In our case, x is the quantum measurement outcome in one-hot notation. A collection of measurement outcomes is a matrix of size $N \times 4$, where N is the number of particles in the chain and 4 is the number of possible outcomes of the tetrahedral IC POVM, which is either [1000], [0100], [0010], or [0001]. h is the external magnetic field. The probability distribution $p[x|z, \theta, h]$ can thus be written as

$$p[x|z, \theta, h] = \prod_{i=1}^N \prod_{j=1}^4 \pi_{ij}(z, h, \theta)^{x_{ij}}, \tag{11}$$

where $\pi_{ij}(z, h, \theta)$ is the neural network in our architecture, and, more precisely, π_{ij} is the probability of the j^{th} outcome of the POVM for the i^{th} spin with $\sum_{j=1}^4 \pi_{ij} = 1$ and $\pi_{ij} \geq 0$. The quantity $p[z]$ is the prior distribution over latent variables, which is simply given by $\mathcal{N}(0, I) = \frac{1}{\sqrt{2\pi}^N} \exp\{-\frac{1}{2}z^T z\}$, with I being the identical covariance matrix. We take the number of latent variables equal to the number of spins, N . Essentially, we want to optimize our VAE so that its probability matches the probability of the quantum measurement outcomes as closely as possible. This can be done using the well-known maximum likelihood estimation:

$$\theta_{MLE} = \operatorname{argmax}_{\theta} \sum_{i=1}^M \log(p[x_i|\theta, h]), \tag{12}$$

where $\{x_i\}_{i=1}^M$ is the data set of outcome measurements. We cannot simply maximize this function using, for example, a gradient descent method, due to the presence of hidden variables in the structure of this function. However, we can overcome this problem by using the Evidence Lower Bound (ELBO) [95] and the reparametrization trick shown in [96]. The detailed description of the procedure is given in the Appendix A.

Once trained, the VAE is a simple and efficient way to produce new samples from its probability distribution. It can be done in three steps. First, we produce a sample from the prior distribution $p[z] = \mathcal{N}(0, I)$. Next, we feed this sample and the external magnetic field value into the neural network decoder $\pi_{ij}(z, \theta, h)$, which returns the matrix of probabilities. Finally, we sample from the matrix of probability $\pi_{ij}(z, \theta, h)$ to generate “fake” outcome measurements. A visual representation of the sampling method is shown in Figure 4.

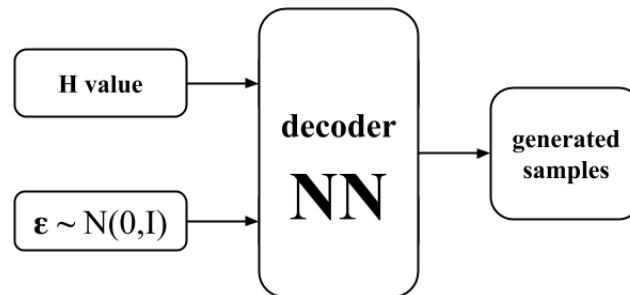


Figure 4. Sampling scheme with the trained variational autoencoder (VAE).

In many problems, gradients of observables with respect to different model parameters yield quantities of interest. For example, one may consider the magnetic differential susceptibility tensor $\chi_{ij} = \partial\mu_i / \partial h_j$. It can be done efficiently by using backpropagation through the VAE architecture but, as samples from the VAE are discrete, a straightforward backpropagation is impossible. In recent papers [97–99], a method called the Gumbel-softmax was introduced to overcome this difficulty through continuous relaxation. The spirit, and therefore the physical meaning of the method, may be understood with a short discussion of the so-called simulated annealing technique, which is often used to solve discrete optimization problems. Broadly speaking, the simulated annealing rests on the introduction of a parameter that acts as an artificial “temperature”, which varies continuously to modify the state of the system in search of a global optimum. Starting from a given state, for some values of the temperature, if the system mostly explores the neighboring states, moving among them and possibly in the vicinity of the “better” ones, i.e., with lower energy, it may get and remain close to a local optimum, or local energy minimum in the thermodynamic language; however, to avoid remaining in a locally optimal region, “bad” moves leading to worse (i.e., higher energy) states are useful to explore the temperature space more completely improving the chance to find a global optimum or at least to be near it. To each move an energy variation, ΔE , is associated; it is the continuous character of the fictitious temperature that makes the discrete problem continuous as the probability $\exp(-\Delta E)/k_B T$ of acceptance of a state is continuous. Although this approach has been known for a long time [100], it remains topical and under active development [101,102]. The method of continuous relaxation we use also exploits such an artificial temperature to make discrete samples continuous.

The Gumbel-softmax trick, consists of three steps:

1. We calculate the matrix of log probabilities, taking element-wise logarithm of decoder network

$$\text{output: } \log \Pi = \begin{bmatrix} \log \pi_{11} & \log \pi_{12} & \dots & \log \pi_{1N} \\ \log \pi_{21} & \log \pi_{22} & \dots & \log \pi_{2N} \\ \log \pi_{31} & \log \pi_{32} & \dots & \log \pi_{3N} \\ \log \pi_{41} & \log \pi_{42} & \dots & \log \pi_{4N} \end{bmatrix},$$

2. We generate a matrix of samples from the standard Gumbel distribution G and sum it up element-wise with the matrix of log probabilities $\log \Pi$: $Z = \log \Pi + G$,
3. Finally, we take the softmax function of the result from the previous step: $x_{\text{soft}}^{\text{fake}}(T) = \text{softmax}(Z/T)$, where T is a temperature of softmax. The softmax function is defined by the expression $\text{softmax}(x_{ij}) = \frac{\exp(x_{ij})}{\sum_i \exp(x_{ij})}$.

The quantity $x_{\text{soft}}^{\text{fake}}(T)$ has a number of remarkable properties: first, it becomes an exact one-hot sample when $T \rightarrow 0$; second, we can backpropagate through soft samples for any $T > 0$. The method is validated in the next section.

Before we proceed to the presentation and discussion of our results, and to better see the added value of the VAE, it is instructive to compare MPS and VAE (NN) in terms of expressibility, i.e., “estimation of MPS states via incomplete local measurements” vs “VAE reconstruction”. As the state of the system is assumed to be unknown, and some measurement outcomes are only known for different magnetic fields, these outcomes are too few for exact tomography. Further, it is known that for a given bond dimension d , the entangled entropy cannot be larger than $\log(d)$; in other words, the bond dimension of MPS places an upper bound on the entangled entropy. Thus, the MPS representation describes well only quantum states with low entangled entropy, i.e., quantum states which satisfy the area law [103,104]. The situation with neural network quantum states (NQS) is different: there is no such a restriction for NQS. Moreover, the existence of NQS with volume-law entanglement [105] shows a promising development of new, and possibly powerful, NN-based approaches to representing many-body quantum systems.

5. Results

Here, we show that the VAE trained on a set of preliminary measurements is capable to describe the physics of the whole family of TFI models. We validate our results by comparing VAE-based calculations with numerically exact calculations performed by variational MPS algorithm [35]. Additionally, to assess the capabilities of the VAE, we consider a spin chain with 32 spins. We calculate the MPS representation of the ground state and extract information from it by performing measurements over the state. The external field in the x -direction is varied from 0 to 2 with a step of 0.1. The VAE is trained on a data set (TFI measurement outcomes) consisting of 10.5 million samples in total: 21 external fields h_x with 500,000 samples per field.

To evaluate the VAE performance, we simply compare directly the numerically exact correlation functions with those reconstructed with our VAE. Those of $n = 1, \dots, 32$, $\langle \sigma_z^1 \sigma_z^n \rangle$, and $\langle \sigma_x^1 \sigma_x^n \rangle$ are shown in Figures 5 and 6, respectively, and we compare the numerically exact and the VAE-based average magnetizations along x , given by $\langle \sigma_x^n \rangle$ for each position of the spin along the chain, in Figure 7. We see that the VAE captures well the physics of the one- and two-point correlation functions. Figure 8 shows the total magnetizations, μ_x and μ_z , in the x and z directions, respectively, with $\mu_i = \frac{1}{N} \sum_{j=1}^N \langle \sigma_i^j \rangle$, and we see that the VAE is a tool well-suited for the description of the quantum phase transition and also finite-size effects: whereas for the infinite TFI chain, i.e., in the thermodynamic limit, the phase transition is observed at $h_x = 1$, and the finite size of the system yields a shift of the critical point at $h_x \approx 0.9$. Also note that in the $T \rightarrow 0$ limit, the magnetization M defined in Equation (3) coincides exactly with the magnetization μ defined above.

A backpropagation algorithm combined with the Gumbel-softmax trick may be used to evaluate the derivative of an output over an input. We use this approach to calculate some elements of a magnetic differential susceptibility tensor $\chi_{ij} = \partial \mu_i / \partial h_j$, in particular, χ_{xx} and χ_{zx} shown in Figure 9. The backpropagation-based magnetic differential susceptibility agrees well with the numerically calculated one (central differences). The main advantage of the backpropagation-based calculation is its numerical efficiency. The VAE may thus be trained with an arbitrary set of external parameters, i.e., not only h_x , but also h_y and h_z , and yield the full differential susceptibility tensor.

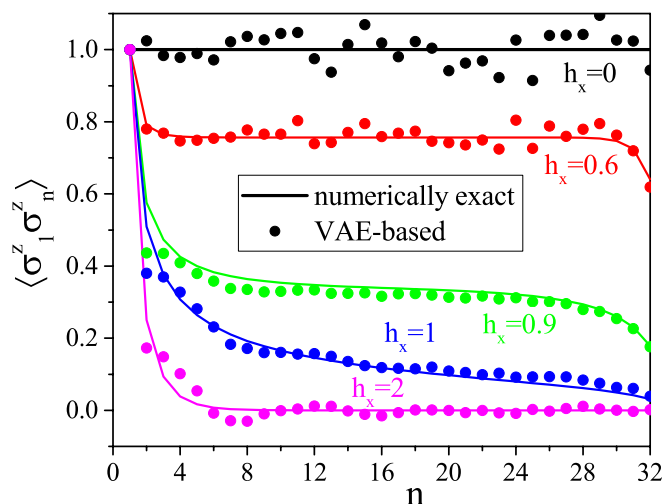


Figure 5. Two-point correlation function $\langle \sigma_1^z \sigma_n^z \rangle$ for different values of external magnetic field h_x .

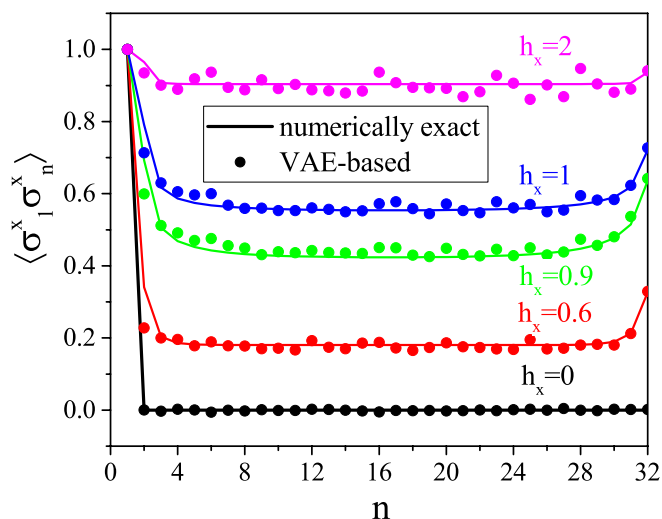


Figure 6. Two-point correlation function $\langle \sigma_1^x \sigma_n^x \rangle$ for different values of external magnetic field h_x .

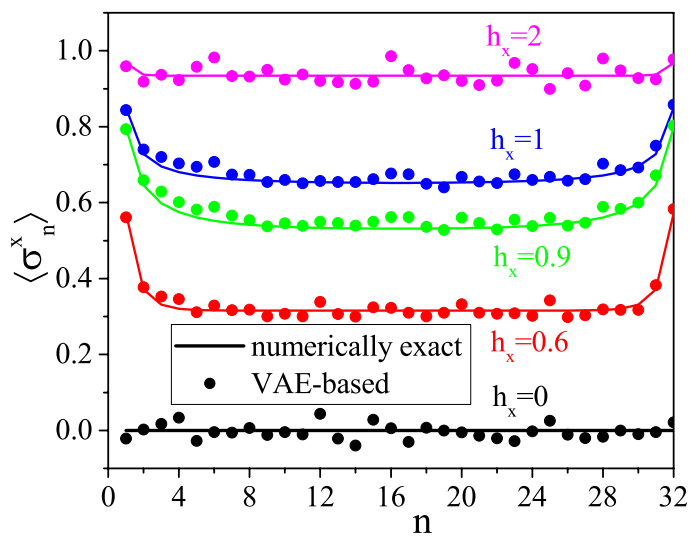


Figure 7. Average magnetization per site along x for different values of external magnetic field h_x .

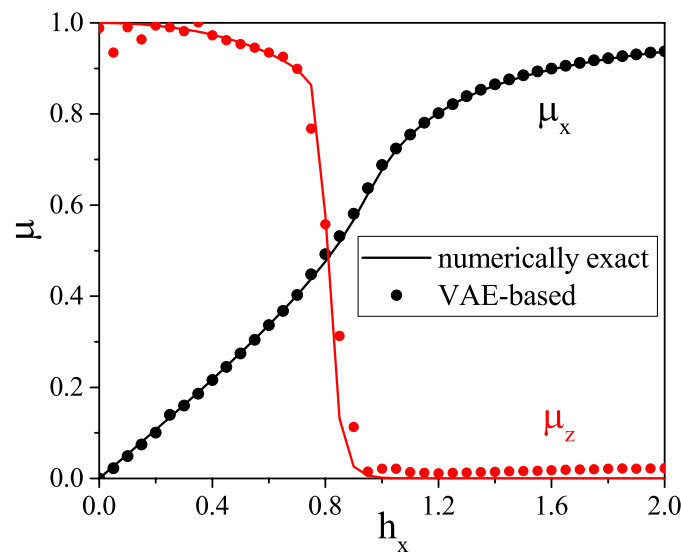


Figure 8. Total magnetization along x and z axes for different values of external magnetic field h_x . The location of the critical region is slightly shifted towards smaller values of h_x due to the finite size of the chain.

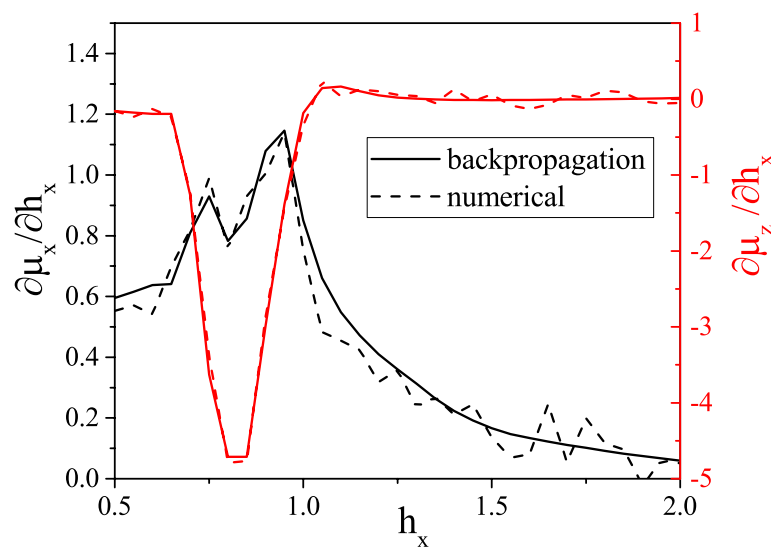


Figure 9. Backpropagation-based and numerical-based (central differences) values of χ_{xx} and χ_{zx} for different values of external magnetic field h_x . Both derivatives slightly fluctuate due to VAE error.

At this stage, we could conclude that the VAE is capable to describe the physics of one- and two-point correlation functions, and therefore the TFI physics. However, notwithstanding the ability of the VAE to yield correlation functions that fit well numerically-exact correlation functions, this is not yet a full proof that it represents quantum states well. To address this point, we consider a small spin chain (five spins with TFI Hamiltonian and an external magnetic field $h_x = 0.9$) for which we calculate both the exact mass function and that estimated from VAE samples. Figure 10 shows that the VAE result again fits the numerically exact mass function with high accuracy. Further, we calculate the Bhattacharyya coefficient [106]: $BC(p_{vae}, p_{exact}) = \sum_{\alpha} p_{exact}[\alpha] \sqrt{\frac{p_{vae}[\alpha]}{p_{exact}[\alpha]}}$ as a function of the external magnetic field h_x . Results reported in Figure 11 show that $BC(p_{vae}, p_{exact}) > 0.99$ over the whole h_x range, which thus proves that the VAE represents a quantum state well, at least for small spin chains.

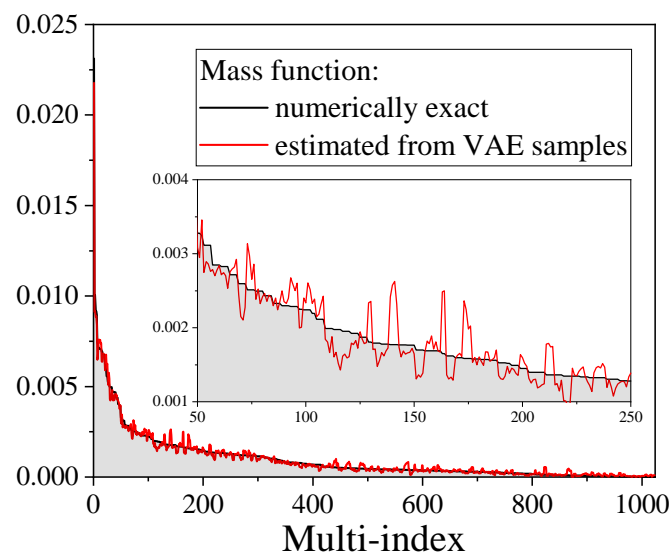


Figure 10. Comparison of two positive-operator valued measure (POVM)-induced mass functions ($P[\alpha] = \text{Tr}(\rho M^\alpha)$) for a chain of size 5: numerically exact mass function and reconstructed from VAE samples mass function. A sequence of indices α has been transformed into a single multi-index. Indices have been ordered to put numerically exact probability in descending order. A good agreement between the mass functions is observed.

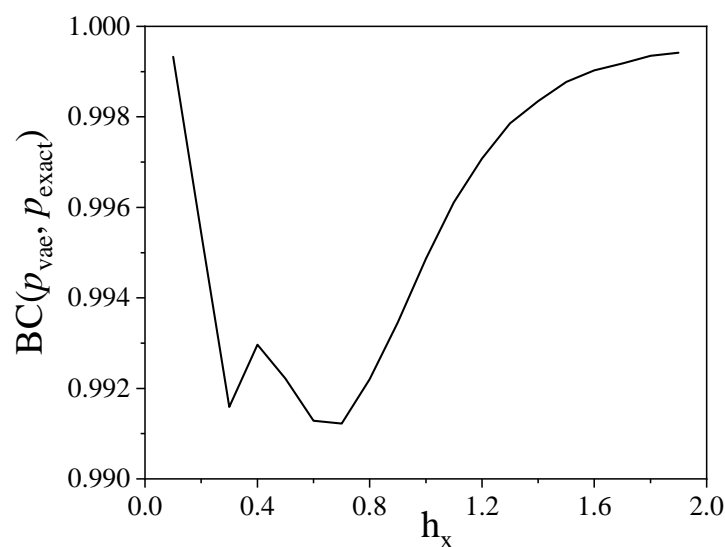


Figure 11. Dependence of the classical fidelity on the external magnetic field. A high predictive accuracy is demonstrated for the whole set of fields.

The structure of the entanglement is another interesting subject that we would like to validate. The essence of entanglement between two parts of the chain, which is split into n left spins and $N - n$ right spins, can be described by the Rényi entropy of the left part of this chain: $S_\alpha = \frac{1}{1-\alpha} \log \text{Tr} \rho_n^\alpha$, where ρ_n is the density matrix of the first n spins in the chain. We estimate the Rényi entropy of order 2: $S_2 = -\log(\text{Tr} \rho^2)$, as it can be efficiently calculated from the matrix product representation of the density matrix and from the VAE samples. However, as sample-based estimation of the entangled entropy has a variance that grows exponentially with the number of spins, we consider a small spin chain of size 10. A direct comparison between the numerically exact and the VAE-based entangled entropies is shown for different values of n in Figure 12. For this particular case, the VAE clearly overestimates the entangled entropy. This undesirable effect is indeed observed for all sizes of spin chains, and even for the spin chain of size 5, for which we have an excellent agreement between the

numerically exact mass function and the VAE-based result. The entropy S_2 is sensitive to small errors in the mass function, but it also appears that the primary method of state reconstruction used in the present work has the following shortcomings.

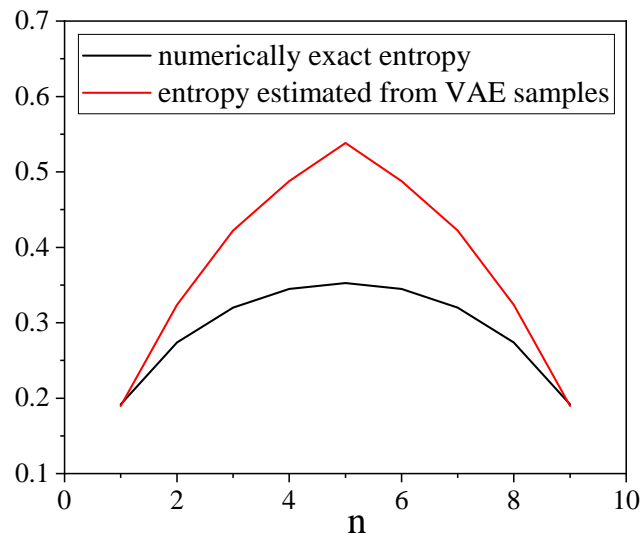


Figure 12. Comparison of the numerically exact Rényi entropy and that reconstructed from the VAE samples for different values of n .

1. If one reconstructs a pure state, the VAE smooths the spectrum of the density matrix and approximates the pure state by a slightly mixed state, as illustrated with a simple example in Figure 13.
2. The VAE does not account the positivity constraints, which yields negative eigenvalues for the density matrix. These negative eigenvalues even appear in the spectrum of the reduced density matrix, as shown in Figure 13.

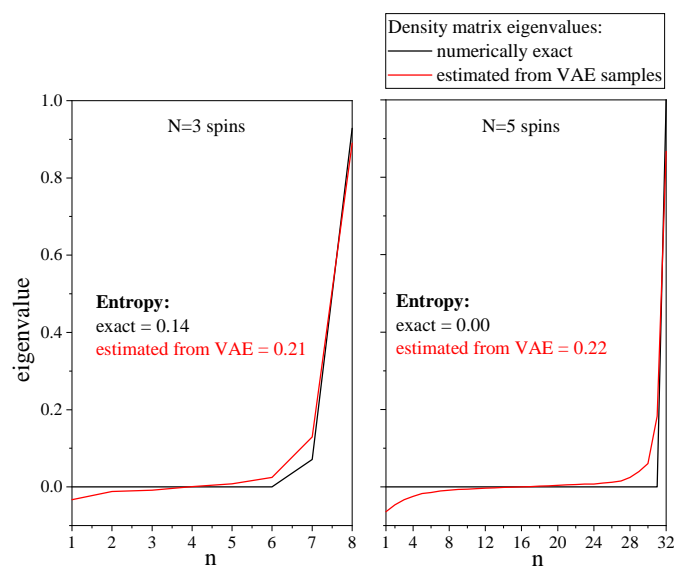


Figure 13. Comparison of numerically exact spectra of density matrices and VAE-estimated spectra. The ground state spectra of the spin chain of size 5 with an external magnetic field $h = 0.9$ is shown on the right panel, and the spectra of the reduced density matrix (last 3 spins) are shown on the left panel.

These drawbacks hinder a robust description of the entanglement structure. In addition to the mismatch between the Rényi entropies (S_2), the entropy of a reduced density matrix can be larger than the entropy of the whole density matrix, which is erroneous. This particular issue, now identified, may be resolved by introduction of a particular regularization term into the VAE loss. This is the object of future work.

Finally, it is also instructive to comment on the memory costs of the use of either MPS or VAE, which is somehow a tricky question, as it is unclear for any NN-based architecture what numbers of layers and neurons per layer are needed because there is no criterion for NN, whereas for the MPS and tensor networks, there is one. Thus, a direct comparison of NN architectures and tensor networks (MPS, etc.) is certainly a difficult task, and in our opinion, likely an impossible one. At this stage, we may say the following. For a given spin chain of size N and maximal entangled entropy between subchains $S = -\text{Tr} \rho \log \rho$, the MPS requires to store approximately $2N \exp(2S)$ complex numbers; this follows from the fact that one then considers N subtensors of size $\exp(S) \times 2 \times \exp(S)$, where $\exp(S)$ is the typical (approximate) size of bond dimension. For a VAE, although it seems that there are no entropic restrictions, the proper quantitative characterization of the “neural network” complexity of a quantum state still is an open question (for tensor networks, it is the entangled entropy). A VAE contains two neural networks: encoder and decoder. To store a feed-forward neural network, one has to store $\sum_i l_{i-1} \times l_i + l_i$ real numbers, with l_i being the number of neurons in the layer number i . In general, one may conclude that the MPS is preferable for low entangled states, and the VAE is preferable for highly entangled states.

6. Conclusions

The thermodynamic study of complex many-body quantum systems still requires the development of new methods, including those that may stem from machine learning. The quantum Ising model, which is of particular importance for practical purposes [107,108], provides a rich framework to test these new methods that are also useful to obtain deeper physical insight into its nonequilibrium dynamics properties such as, e.g., quantum fluctuations propagation [109]. In the present work, we studied the ability of a VAE to reconstruct the physics of quantum many-body systems, using the transverse-field Ising model as a nontrivial example. We used the IC POVM to map the quantum problem onto a probabilistic domain and vice versa. We trained the VAE on a set of samples from the transformed quantum problem, and our numerical experiments show the following results.

- For a large system (32 spins), the VAE’s reliability is verified by comparing one- and two-point correlation functions.
- For small system (five spins), the VAE’s reliability is verified by direct comparison of mass functions.
- The VAE can capture a quantum phase transition.
- The response functions (magnetic differential susceptibility tensor) can be obtained using backpropagation through VAE.
- Despite the very good agreement between the VAE-based mass function and the true mass function, the VAE shows limited performance with the determination of the entangled entropy. This point is the object of further development.

Our method can be extended to any other thermodynamic system by introduction of the temperature as an external parameter, thereby considering also thermal phase transitions. As one can calculate different thermodynamic quantities by applying backpropagation through VAE, a worthwhile and highly complex system to study would be water under its difference phases, so as to test recent new ideas and models [110,111].

Our code for our numerical experiments is available on the GitHub repository website [112].

Author Contributions: Conceptualization, I.A.L., S.N.F., and H.O.; methodology, I.A.L. and A.R.; software, I.A.L., A.R., and P.J.S.; validation, all authors; writing-original draft preparation, I.A.L., A.R., P.J.S., and H.O.; writing-review and editing, S.N.F. and H.O.

Funding: This research was supported by the Russian Foundation for Basic Research grants under the Project No. 18-37-00282 and the Project No. 18-37-20073. This research was also partially supported by the Skoltech NGP Program (Skoltech-MIT joint project).

Acknowledgments: The authors thank Stepan Vintskevich for fruitful discussions. The authors also thank Google Colaboratory for providing access to GPU for the acceleration of computations.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript.

VAE	Variational Autoencoder
MPS	Matrix product state
TFI	Transverse-field Ising
IC	Informationally incomplete
POVM	Positive-operator valued measure
ELBO	Evidence lower bound
NN	Neural network
KL	Kullback–Leibler
DMRG	Density matrix renormalization group

Appendix A. VAE: Training and Implementation Details

When training our VAE, we find the arg maximum of the logarithmic likelihood $\mathcal{L}(\theta)$ w.r.t. its parameters θ :

$$\theta_{\text{MLE}} = \underset{\theta}{\operatorname{argmax}} \mathcal{L}(\theta) = \underset{\theta}{\operatorname{argmax}} \log(p[x|\theta, h]), \quad (\text{A1})$$

Equation (A1) cannot directly be evaluated, because of hidden variables in the structure of $p[x|\theta, h]$. We can, however, simplify this problem by introducing a distribution over hidden variables z . Remember that the probability distribution can be described as $p[x|\theta, h] = \int p[x|z, \theta, h]p[z]dz$, so that the expression for the log likelihood becomes

$$\mathcal{L}(\theta) = \log \left(\int p[x|z, \theta, h]p[z]dz \right). \quad (\text{A2})$$

We can then use a mathematical trick that might seem counterintuitive at first glance, but ultimately becomes quite powerful. We multiply the function inside the integral by $\frac{q[z|x, \tilde{\theta}, h]}{q[z|x, \tilde{\theta}, h]} = 1$, where $q[z|x, \tilde{\theta}, h]$ is some arbitrary distribution that can be adjusted with $\tilde{\theta}$, so that

$$\begin{aligned} \mathcal{L}(\theta) &= \log \left(\int p[x|z, \theta, h]p[z]dz \right) = \log \left(\int \frac{q[z|x, \tilde{\theta}, h]}{q[z|x, \tilde{\theta}, h]} p[x|z, \theta, h]p[z]dz \right) \\ &= \log \left(\mathbb{E}_{q[z|x, \tilde{\theta}, h]} p[x|z, \theta, h] \frac{p[z]}{q[z|x, \tilde{\theta}, h]} \right) \end{aligned} \quad (\text{A3})$$

where the quantity $\mathbb{E}_{f[x]}$ denotes the expectation value w.r.t some distribution $f[x]$. We can then use Jensen's inequality to show that

$$\log \left(\mathbb{E}_{q[z|x, \tilde{\theta}, h]} p[x|z, \theta, h] \frac{p[z]}{q[z|x, \tilde{\theta}, h]} \right) \geq \mathbb{E}_{q[z|x, \tilde{\theta}, h]} \log \left(p[x|z, \theta, h] \frac{p[z]}{q[z|x, \tilde{\theta}, h]} \right). \quad (\text{A4})$$

where the rhs of this inequality is the lower bound of the log likelihood, as it will always be greater than or equal to the lower bound, and equality can always be achieved by a proper choice of q if it is in a complex enough family.

Maximizing the lower bound is equivalent to maximizing the log likelihood. We can decompose this lower bound term into two terms:

$$\mathcal{L}(\theta) \geq \text{ELBO}(\theta, \tilde{\theta}) = \mathbb{E}_{q[z|x, \tilde{\theta}, h]} \log(p[x|z, \theta, h]) - \int q[z|x, \tilde{\theta}, h] \log \frac{q[z|x, \tilde{\theta}, h]}{p[z]} dz \tag{A5}$$

Note that the second term is equivalent to the Kullback–Leibler divergence $KL(q[z|x, \tilde{\theta}, h] || p[z])$. In our case, we picked the particular distribution forms that reflect the structure of our problem:

$$\begin{aligned} p[x|z, \theta, h] &= \prod_{i=1}^N \prod_{j=1}^4 \pi_{ij}(z, \theta, h)^{x_{ij}}, \\ q[z|x, \tilde{\theta}, h] &= \mathcal{N}(\mu_i(x, \tilde{\theta}, h), \text{Diag}(\sigma_i^2(x, \tilde{\theta}, h))), \\ P[z] &= \mathcal{N}(0, I) \end{aligned} \tag{A6}$$

where μ_i and σ_i are given by the encoder neural network, and π_{ij} is given by the decoder neural network, with $\sum_{j=1}^4 \pi_{ij} = 1$ and $\pi_{ij} \geq 0$, which can be achieved by applying the softmax function to the output of the neural network. Now, we can use the reparametrization trick to change the variable in the integral $z = \sigma_j(x, \tilde{\theta}, h)\varepsilon + \mu_j(x, \tilde{\theta}, h)$, where $\varepsilon_j \sim \mathcal{N}(0, I)$, to simplify this expression to

$$\begin{aligned} \text{ELBO}(\theta, \tilde{\theta}) &= \sum_{i=1}^N \sum_{j=1}^4 x_{ij} \langle \log(\pi_{ij}(\sigma_i(x, \tilde{\theta}, h)\varepsilon + \mu_i(x, \tilde{\theta}, h), \theta, h)) \rangle_{\varepsilon_j \sim \mathcal{N}(0, I)} \\ &\quad - \sum_{i=1}^N \left(\log \sigma_i(x, \tilde{\theta}, h) - \frac{\sigma_i^2(x, \tilde{\theta}, h) + \mu_i^2(x, \tilde{\theta}, h) - 1}{2} \right). \end{aligned} \tag{A7}$$

The first term is the cross-entropy, which pushes the probability distribution to be as close as possible to the data. The second term is the regularizer, which forces the latent variable z not to diverge too much from the normal distribution $\mathcal{N}(0, I)$, so that the VAE can be used to generate new data once it is trained. Note that both x_{ij} and σ_i must be positive. Instead of adding a constraint to the VAE, which would be difficult to do, we train the VAE for the variables $\Pi = \log \pi$ and $\xi = 2 \log \sigma$. Equation (A7) then becomes

$$\begin{aligned} \text{ELBO}(\theta, \tilde{\theta}) &= \sum_{i=1}^N \sum_{j=1}^4 x_{ij} \langle \Pi_{ij}(e^{\xi_i(x, \tilde{\theta}, h)/2} \varepsilon + \mu_i(x, \tilde{\theta}, h), \theta, h) \rangle_{\varepsilon_j \sim \mathcal{N}(0, I)} \\ &\quad - \frac{1}{2} \sum_{i=1}^N \left(\xi_i(x, \tilde{\theta}, h) - e^{\xi_i(x, \tilde{\theta}, h)} - \mu_i^2(x, \tilde{\theta}, h) + 1 \right). \end{aligned} \tag{A8}$$

Now, $\text{ELBO}(\theta, \tilde{\theta})$ can be effectively optimized using gradient descent methods, averaging over ε can be done by sampling. Generalizing to a data set of size M : $\{x^k\}_{k=1}^M$ can be easily done and is shown by

$$\begin{aligned} \text{ELBO}(\theta, \tilde{\theta}) &= \sum_{k=1}^M \sum_{i=1}^N \sum_{j=1}^4 x_{ij}^k \langle \Pi_{ij}(e^{\xi_i(x^k, \tilde{\theta}, h)/2} \varepsilon + \mu_i(x^k, \tilde{\theta}, h), \theta, h) \rangle_{\varepsilon_j \sim \mathcal{N}(0, I)} \\ &\quad - \frac{1}{2} \sum_{k=1}^M \sum_{i=1}^N \left(\xi_i(x^k, \tilde{\theta}, h) - e^{\xi_i(x^k, \tilde{\theta}, h)} - \mu_i^2(x^k, \tilde{\theta}, h) + 1 \right). \end{aligned} \tag{A9}$$

A visual representation of the VAE architecture is shown in Figure A1.

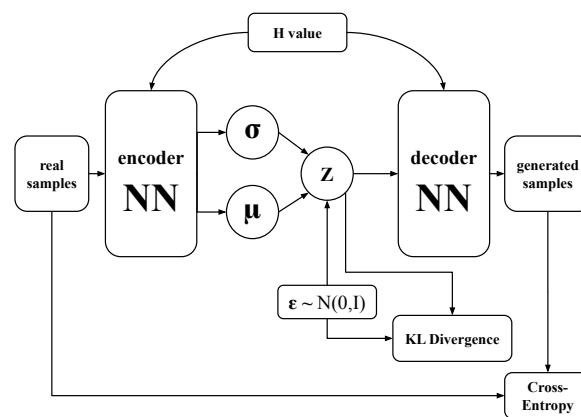


Figure A1. Architecture of the variational autoencoder.

To solve the optimization problem, we use Adam optimizer [113] with standard parameters ($\text{lr} = 0.001, \beta_1 = 0.9, \beta_2 = 0.999$). For the encoder and decoder, we use fully-connected neural networks with two hidden layers and 256 neurons on each. We train the VAE using batches of size 100,000 samples and for 750 epochs.

Appendix B. Sampling from POVM-Induced Mass Function

The mass function induced by POVM $P[\alpha_1, \alpha_2, \dots, \alpha_N]$ has a form of matrix product state. Thus, one can easily calculate any marginal mass function because a summation over any α can be done locally. Any conditional mass functions can be also calculated by using marginal mass functions. Thus, one can calculate chain decomposition of the whole mass function:

$$P[\alpha_1, \alpha_2, \dots, \alpha_N] = P[\alpha_N]P[\alpha_{N-1}|\alpha_N]P[\alpha_{N-2}|\alpha_{N-1}, \alpha_N] \dots P[\alpha_1|\alpha_2, \dots, \alpha_N] \quad (\text{A10})$$

With this decomposition, one can produce a sample $\tilde{\alpha}_N$ from $P[\alpha_N]$ first, then a sample $\tilde{\alpha}_{N-1}$ from $P[\alpha_{N-1}|\tilde{\alpha}_N]$, and continue up to the end of the chain. The obtained set $\{\tilde{\alpha}_1, \tilde{\alpha}_2, \dots, \tilde{\alpha}_N\}$ is a valid sample from the mass function.

References

1. Muller, I. *A History of Thermodynamics*; Springer: Berlin/Heidelberg, Germany, 2007.
2. Onsager, L. Reciprocal Relations in Irreversible Processes. I. *Phys. Rev.* **1931**, *37*, 405–426. [[CrossRef](#)]
3. De Groot, S.R. *Thermodynamics of Irreversible Processes*; Interscience: New York, NY, USA, 1958.
4. Le Bellac, M.; Mortessagne, F.; Batrouni, G.G. *Equilibrium and Non-Equilibrium Statistical Thermodynamics*; Cambridge University Press: Cambridge, UK, 2006.
5. Apertet, Y.; Ouerdane, H.; Goupil, C.; Lecoeur, P. Revisiting Feynman's ratchet with thermoelectric transport theory. *Phys. Rev. E* **2014**, *90*, 012113. [[CrossRef](#)] [[PubMed](#)]
6. Goupil, C.; Ouerdane, H.; Herbert, E.; D'Angelo, Y.; Lecoeur, P. Closed-loop approach to thermodynamics. *Phys. Rev. E* **2016**, *94*, 032136. [[CrossRef](#)]
7. Andresen, B. Current trends in finite-time thermodynamics. *Angew. Chem.-Int. Edit.* **2011**, *50*, 2690–2704. [[CrossRef](#)] [[PubMed](#)]
8. Ouerdane, H.; Apertet, Y.; Goupil, C.; Lecoeur, P. Continuity and boundary conditions in thermodynamics: From Carnot's efficiency to efficiencies at maximum power. *Eur. Phys. J. Spec. Top.* **2015**, *224*, 839–864. [[CrossRef](#)]
9. Apertet, Y.; Ouerdane, H.; Goupil, C.; Lecoeur, P. True nature of the Curzon-Ahlborn efficiency. *Phys. Rev. E* **2017**, *96*, 022119. [[CrossRef](#)]
10. Boltzmann, L. Über die beziehung dem zweiten Haubtsatze der mechanischen Warmetheorie und der Wahrscheinlichkeitsrechnung respektive den Satzen über das Warmegleichgewicht. *Wiener Berichte* **1877**, *76*, 373–435.

11. Gibbs, J.W. *Elementary Principles in Statistical Mechanics*; Charles Scribner's Sons: New York, NY, USA, 1902.
12. Penrose, O. Foundations of statistical mechanics. *Rep. Prog. Phys.* **1979**, *42*, 1937–2006. [[CrossRef](#)]
13. Goldstein, S.; Lebowitz, J.L.; Zanghì, N. Gibbs and Boltzmann entropy in classical and quantum mechanics. *arXiv* **2019**, arXiv:1903.11870. Available online: <https://arxiv.org/abs/1903.11870> (accessed on 6 November 2019).
14. Shannon, C.E. A mathematical theory of communication. *Bell Labs Tech. J.* **1948**, *27*, 379–423. [[CrossRef](#)]
15. von Neumann, J. *Mathematical Foundations of Quantum Mechanics. New Edition*; Princeton University Press: Princeton, NJ, USA, 2018.
16. Datta, S. *Electronic Transport in Mesoscopic Systems*; Cambridge University Press: Cambridge, UK, 1995.
17. Heikkilä, T.T. *The Physics of Nanoelectronics*; Oxford University Press: Oxford, UK, 2013.
18. Chomaz, P.; Colonna, M.; Randrup, J. Nuclear spinodal fragmentation. *Phys. Rep.* **2004**, *389*, 263–440. [[CrossRef](#)]
19. Bressanini, D.; Morosi, G.; Mella, M. Robust wave function optimization procedures in quantum Monte Carlo methods. *J. Chem. Phys.* **2002**, *116*, 5345–5350. [[CrossRef](#)]
20. Feiguin, A.E.; White, S.R. Finite-temperature density matrix renormalization using an enlarged Hilbert space. *Phys. Rev. B* **2005**, *72*, 220401. [[CrossRef](#)]
21. Deutsch, J.M. Quantum statistical mechanics in a closed system. *Phys. Rev. A* **1991**, *43*, 2046–2049. [[CrossRef](#)]
22. Srednicki, M. Chaos and quantum thermalization. *Phys. Rev. E* **1994**, *50*, 888–901. [[CrossRef](#)]
23. Rigol, M.; Dunjko, V.; Olshanii, M. Thermalization and its mechanism for generic isolated quantum systems. *Nature* **2008**, *452*, 854–858. [[CrossRef](#)]
24. Dymarsky, A.; Lashkari, N.; Liu, H. Subsystem eigenstate thermalization hypothesis *Phys. Rev. E* **2018**, *97*, 012140. [[CrossRef](#)] [[PubMed](#)]
25. Dymarsky, A. Mechanism of macroscopic equilibration of isolated quantum systems *Phys. Rev. B* **2019**, *99*, 224302. [[CrossRef](#)]
26. Carleo, G.; Becca, F.; Schiró, M.; Fabrizio, M. Localization and glassy dynamics of many-body quantum systems. *Sci. Rep.* **2012**, *2*, 243. [[CrossRef](#)] [[PubMed](#)]
27. Chen, L.; Gelin, M.; Zhao, Y. Dynamics of the spin-boson model: A comparison of the multiple Davydov D_1 , $D_{1.5}$, D_2 Ansätze. *Chem. Phys.* **2018**, *515*, 108–118. [[CrossRef](#)]
28. Lanyon, B.; Maier, C.; Holzäpfel, M.; Baumgratz, T.; Hempel, C.; Jurcevic, P.; Dhand, I.; Buyskikh, A.; Daley, A.; Cramer, M.; et al. Efficient tomography of a quantum many-body system. *Nat. Phys.* **2017**, *13*, 1158. [[CrossRef](#)]
29. Liao, H.J.; Liu, J.G.; Wang, L.; Xiang, T. Differentiable programming tensor networks. *Phys. Rev. X* **2019**, *9*, 031041. [[CrossRef](#)]
30. Fetter, A.L.; Walecka, J.D. *Quantum Theory of Many-Particle Systems*; Dover: New York, NY, USA, 2003.
31. Frésard, R.; Kroha, J.; Wölfle, P. The pseudoparticle approach to strongly correlated electron systems. In *Strongly Correlated Systems*; Avella, A., Mancini, F., Eds.; Springer: Berlin/Heidelberg, Germany, 2011; Volume 171.
32. Georges, A.; Kotliar, G.; Krauth, W.; Rozenberg, M.J. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.* **1996**, *68*, 13–125. [[CrossRef](#)]
33. Negele, J.W.; Orland, H. *Quantum Many-Particle Systems*; Perseus Books: New York, NY, USA, 1998.
34. Foulkes, W.; Mitas, L.; Needs, R.; Rajagopal, G. Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.* **2001**, *73*, 33. [[CrossRef](#)]
35. Orús, R. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Ann. Phys.* **2014**, *349*, 117–158. [[CrossRef](#)]
36. Orús, R. Tensor networks for complex quantum systems. *Nat. Rev. Phys.* **2019**, *1*, 538–550. [[CrossRef](#)]
37. Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **2011**, *326*, 96–192. [[CrossRef](#)]
38. Vidal, G. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* **2003**, *91*, 147902. [[CrossRef](#)]
39. Evenbly, G.; Vidal, G. Quantum criticality with the multiscale entanglement renormalization ansatz. In *Strongly Correlated Systems*; Springer: Berlin/Heidelberg, Germany, 2013; pp. 99–130.
40. Pollock, F.A.; Rodríguez-Rosario, C.; Frauenheim, T.; Paternostro, M.; Modi, K. Non-Markovian quantum processes: Complete framework and efficient characterization. *Phys. Rev. A* **2018**, *97*, 012127. [[CrossRef](#)]

41. Luchnikov, I.; Vintskevich, S.; Ouerdane, H.; Filippov, S. Simulation complexity of open quantum dynamics: Connection with tensor networks. *Phys. Rev. Lett.* **2019**, *122*, 160401. [CrossRef]
42. Taranto, P.; Pollock, F.A.; Modi, K. Memory strength and recoverability of non-Markovian quantum stochastic processes. *arXiv* **2019**, arXiv:1907.12583. Available online: <https://arxiv.org/abs/1907.12583> (accessed on 6 November 2019).
43. Milz, S.; Pollock, F.A.; Modi, K. Reconstructing non-Markovian quantum dynamics with limited control. *Phys. Rev. A* **2018**, *98*, 012108. [CrossRef]
44. Luchnikov, I.A.; Vintskevich, S.V.; Grigoriev, D.A.; Filippov, S.N. Machine learning of Markovian embedding for non-Markovian quantum dynamics. *arXiv* **2019**, arXiv:1902.07019. Available online: <https://arxiv.org/abs/1902.07019> (accessed on 6 November 2019).
45. Verstraete, F.; Murg, V.; Cirac, J.I. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.* **2008**, *57*, 143–224. [CrossRef]
46. Levin, M.; Nave, C.P. Tensor renormalization group approach to two-dimensional classical lattice models. *Phys. Rev. Lett.* **2007**, *99*, 120601. [CrossRef]
47. Evenbly, G.; Vidal, G. Tensor network renormalization. *Phys. Rev. Lett.* **2015**, *115*, 180405. [CrossRef] [PubMed]
48. Gemmer, J.; Michel, M. *Quantum Thermodynamics*; Springer: Berlin/Heidelberg, Germany, 2009.
49. Kosloff, R. Quantum thermodynamics and open-systems modeling. *J. Phys. Chem.* **2019**, *150*, 204105. [CrossRef]
50. Allahverdyan, A.E.; Johal, R.S.; Mahler, G. Work extremum principle: Structure and function of quantum heat engines. *Phys. Rev. E* **2008**, *77*, 041118. [CrossRef]
51. Thomas, G.; Johal, R.S. Coupled quantum Otto cycle. *Phys. Rev. E* **2011**, *83*, 031135. [CrossRef]
52. Makhlin, Y.; Schön, G.; Shnirman, A. Quantum-state engineering with Josephson-junction devices. *Rev. Mod. Phys.* **2001**, *73*, 357–400. [CrossRef]
53. Navez, P.; Sowa, A.; Zagoskin, A. Entangling continuous variables with a qubit array. *Phys. Rev. B* **2019**, *100*, 144506. [CrossRef]
54. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2006.
55. Turing, M.A. Computing machinery and intelligence. *Mind* **1950**, *59*, 433–460. [CrossRef]
56. Crevier, D. *AI: The Tumultuous Search for Artificial Intelligence*; BasicBooks: New York, NY, USA, 1993.
57. Biamonte, J.; Wittek, P.; Pancotti, N.; Rebentrost, P.; Wiebe, N.; Lloyd, S. Quantum machine learning. *Nature* **2017**, *549*, 195–202. [CrossRef] [PubMed]
58. Carleo, G.; Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **2017**, *355*, 602–606. [CrossRef] [PubMed]
59. Torlai, G.; Mazzola, G.; Carrasquilla, J.; Troyer, M.; Melko, R.; Carleo, G. Neural-network quantum state tomography. *Nat. Phys.* **2018**, *14*, 447. [CrossRef]
60. Tiunov, E.S.; Tiunova, V.V.; Ulanov, A.E.; Lvovsky, A.I.; Fedorov, A.K. Experimental quantum homodyne tomography via machine learning. *arXiv* **2019**, arXiv:1907.06589. Available online: <https://arxiv.org/abs/1907.06589> (accessed on 6 November 2019).
61. Choo, K.; Neupert, T.; Carleo, G. Study of the two-dimensional frustrated J1-J2 model with neural network quantum states. *Phys. Rev. B* **2019**, *100*, 124125. [CrossRef]
62. Sharir, O.; Levine, Y.; Wies, N.; Carleo, G.; Shashua, A. Deep autoregressive models for the efficient variational simulation of many-body quantum systems. *arXiv* **2019**, arXiv:1902.04057. Available online: <https://arxiv.org/abs/1902.04057> (accessed on 6 November 2019).
63. Wu, D.; Wang, L.; Zhang, P. Solving statistical mechanics using variational autoregressive networks. *Phys. Rev. Lett.* **2019**, *122*, 080602. [CrossRef]
64. Kharkov, Y.A.; Sotskov, V.E.; Karazeev, A.A.; Kiktenko, E.O.; Fedorov, A.K. Revealing quantum chaos with machine learning. *arXiv* **2019**, arXiv:1902.09216. Available online: <https://arxiv.org/abs/1902.09216> (accessed on 6 November 2019).
65. Rocchetto, A.; Grant, E.; Strelchuk, S.; Carleo, G.; Severini, S. Learning hard quantum distributions with variational autoencoders. *npj Quantum Inf.* **2018**, *4*, 28. [CrossRef]
66. Carrasquilla, J.; Torlai, G.; Melko, R.G.; Aolita, L. Reconstructing quantum states with generative models. *Nat. Mach. Intell.* **2019**, *1*, 155. [CrossRef]

67. Generative Models for Physicists. Lecture note. Available online: <http://wangleiphy.github.io/lectures/PILTutorial.pdf> (accessed on 7 November 2019).
68. Hewson, A.C. *The Kondo Problem to Heavy Fermions*; Cambridge University Press: Cambridge, UK, 1993.
69. Coleman, P. Heavy fermions: Electrons at the edge of magnetism. In *Handbook of Magnetism and Advanced Magnetic Materials*; Kronmüller, H., Parkin, S., Eds.; John Wiley & Sons: Chichester, UK, 2007.
70. Sachdev, S. *Quantum Phase Transitions*; Cambridge University Press: Cambridge, UK, 2000.
71. Coleman, P.; Schofield, A. Quantum criticality. *Nature* **2000**, *433*, 226–229. [[CrossRef](#)] [[PubMed](#)]
72. Anderson, P.W. Localized Magnetic States in Metals. *Phys. Rev.* **1961**, *124*, 41–53. [[CrossRef](#)]
73. Frésard, R.; Ouerdane, H.; Kopp, T. Slave bosons in radial gauge: a bridge between path integral and Hamiltonian language. *Nucl. Phys. B* **2007**, *785*, 286–306. [[CrossRef](#)]
74. Frésard, R.; Ouerdane, H.; Kopp, T. Barnes slave-boson approach to the two-site single-impurity Anderson model with non-local interaction. *EPL* **2008**, *82*, 31001. [[CrossRef](#)]
75. Diu, B.; Guthmann, C.; Lederer, D.; Roulet, B. *Physique Statistique*; Éditions Hermann: Paris, France, 1996.
76. Mila, F. Frustrated spin systems. In *Many-Body Physics: From Kondo to Hubbard*; Pavarini, E., Koch, E., Coleman, P., Eds.; Verlag des Forschungszentrum Jülich: Kreis Düren, Rheinland, 2015.
77. Refael, G.; Moore, J.E. Entanglement Entropy of Random Quantum Critical Points in One Dimension. *Phys. Rev. Lett.* **2004**, *93*. [[CrossRef](#)]
78. Schollwóck, U. The density-matrix renormalization group. *Rev. Mod. Phys.* **2005**, *77*, 259–315. [[CrossRef](#)]
79. Ising, E. Beitrag zur Theorie des Ferromagnetismus. *Z. Phys.* **1925**, *31*, 253–258. [[CrossRef](#)]
80. Kramers, H.A.; Wannier, G.H. Statistics of the two-dimensional ferromagnet. Part I. *Phys. Rev.* **1941**, *60*, 252–262. [[CrossRef](#)]
81. Ovchinnikov, A.A.; Dmitriev, D.V.; Krivnov, V.Y.; Cheranovskii, V.O. Antiferromagnetic Ising chain in a mixed transverse and longitudinal magnetic field. *Phys. Rev. B* **2003**, *68*, 214406. [[CrossRef](#)]
82. Coldea, R.; Tennant, D.A.; Wheeler, E.M.; Wawrzynska, E.; Prabhakaran, D.; Telling, M.; Habicht, K.; Smeibidl, P.; K, K. Quantum criticality in an Ising chain: Experimental evidence for emergent E_8 symmetry. *Science* **2010**, *327*, 177–180. [[CrossRef](#)] [[PubMed](#)]
83. Sachdev, S.; Keimer, B. Quantum criticality. *Phys. Today* **2011**, *64*, 29–35. [[CrossRef](#)]
84. Matsubara, T. A new approach to quantum statistical mechanics. *Prog. Theor. Exp.* **1955**, *14*, 351–378. [[CrossRef](#)]
85. Kogut, J.B. An introduction to lattice gauge theory and spin systems. *Rev. Mod. Phys.* **1979**, *51*, 659–713. [[CrossRef](#)]
86. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In Proceedings of the NIPS: Advances in Neural Information Processing Systems 25, Stateline, NV, USA, 3–8 December 2012; pp. 1097–1105.
87. Holevo, A.S. *Probabilistic and Statistical Aspects of Quantum Theory*; Springer: Berlin/Heidelberg, Germany, 2011; Volume 1.
88. Filippov, S.N.; Man'ko, V.I. Inverse spin-s portrait and representation of qudit states by single probability vectors. *J. Russ. Laser Res.* **2010**, *31*, 32–54. [[CrossRef](#)]
89. Appleby, M.; Fuchs, C.A.; Stacey, B.C.; Zhu, H. Introducing the Qplex: A novel arena for quantum theory. *Eur. Phys. J. D* **2017**, *71*, 197. [[CrossRef](#)]
90. Caves, C.M. Symmetric informationally complete POVMs - UNM Information Physics Group internal report (1999). Available online: <http://info.phys.unm.edu/~caves/reports/infopovm.pdf> (accessed on 7 November 2019).
91. Myung, I.J. Tutorial on maximum likelihood estimation. *J. Math. Psychol.* **2003**, *47*, 90–100. [[CrossRef](#)]
92. Filippov, S.N.; Man'ko, V.I. Symmetric informationally complete positive operator valued measure and probability representation of quantum mechanics. *J. Russ. Laser Res.* **2010**, *31*, 211–231. [[CrossRef](#)]
93. mpnum: A Matrix Product Representation Library for Python. Available online: <https://mpnum.readthedocs.io/en/latest/> (accessed on 7 November 2019).
94. Sohn, K.; Lee, H.; Yan, X. Learning structured output representation using deep conditional generative models. In Proceedings of the NIPS: Advances in Neural Information Processing Systems 28, Montreal, QC, Canada, 7–12 December 2015; pp. 3483–3491.
95. Kingma, D.P.; Welling, M. Auto-encoding variational Bayes. *arXiv* **2013**, arXiv:1312.6114. Available online: <https://arxiv.org/abs/1312.6114> (accessed on 6 November 2019).

96. Rezende, D.J.; Mohamed, S.; Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In Proceedings of the 31st International Conference on Machine Learning (ICML), Beijing, China, 21–26 June 2014; Volume 32.
97. Jang, E.; Gu, S.; Poole, B. Categorical reparameterization with Gumbel-softmax. *arXiv* **2016**, arXiv:1611.01144. Available online: <https://arxiv.org/abs/1611.01144> (accessed on 6 November 2019).
98. Kusner, M.J.; Hernández-Lobato, J.M. Gans for sequences of discrete elements with the Gumbel-softmax distribution. *arXiv* **2016**, arXiv:1611.04051. Available online: <https://arxiv.org/abs/1611.04051> (accessed on 6 November 2019).
99. Maddison, C.J.; Mnih, A.; Teh, Y.W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv* **2016**, arXiv:1611.00712. Available online: <https://arxiv.org/abs/1611.00712> (accessed on 6 November 2019).
100. Metropolis, N.; Rosenbluth, A.W.; Rosenbluth, M.N.; Teller, A.H.; Teller, E. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **1953**, *21*, 1087–1092. [[CrossRef](#)]
101. Das, A.; Chakrabarti, B.K. Colloquium: Quantum annealing and analog quantum computation. *Rev. Mod. Phys.* **2008**, *80*, 1061–1081. [[CrossRef](#)]
102. Yavorsky, A.; Markovich, L.A.; Polyakov, E.A.; Rubtsov, A.N. Highly parallel algorithm for the Ising ground state searching problem. *arXiv* **2019**, arXiv:1907.05124. Available online: <https://arxiv.org/abs/1907.05124> (accessed on 6 November 2019).
103. Verstraete, F.; Cirac, J.I. Matrix product states represent ground states faithfully. *Phys. Rev. B* **2006**, *73*, 094423. [[CrossRef](#)]
104. Eisert, J.; Cramer, M.; Plenio, M.B. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.* **2010**, *82*, 277–306. [[CrossRef](#)]
105. Deng, D.-L.; Li, X.; Das Sarma, S. Quantum entanglement in neural network states. *Phys. Rev. X* **2017**, *7*, 021021. [[CrossRef](#)]
106. Bhattacharyya, A. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.* **1943**, *35*, 99–109.
107. Boixo, S.; Ronnow, T.F.; Isakov, S.V.; Wang, Z.; Wecker, D.; Lidar, D.A.; Martinis, J.M.; Troyer, M. Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.* **2014**, *10*, 218–224. [[CrossRef](#)]
108. Denchev, V.S.; Boixo, S.; Isakov, S.V.; Ding, N.; Babbush, R.; Smelyanskiy, V.; Martinis, J.; Neven, H. What is the computational value of finite-range tunneling? *Phys. Rev. X* **2016**, *6*, 031015. [[CrossRef](#)]
109. Navez, P.; Tsironis, G.P.; Zagoskin, A.M. Propagation of fluctuations in the quantum Ising model. *Phys. Rev. B* **2017**, *95*, 064304. [[CrossRef](#)]
110. Volkov, A.A.; Artemov, V.G.; Pronin, A.V. A radically new suggestion about the electrodynamics of water: Can the pH index and the Debye relaxation be of a common origin? *EPL* **2014**, *106*, 46004. [[CrossRef](#)]
111. Artemov, V.G. A unified mechanism for ice and water electrical conductivity from direct current to terahertz. *Phys. Chem. Chem. Phys.* **2019**, *21*, 8067–8072. [[CrossRef](#)] [[PubMed](#)]
112. Github Repository with Code. Available online: <https://github.com/LuchnikovI/Representation-of-quantum-many-body-states-via-VAE> (accessed on 7 November 2019).
113. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980. Available online: <https://arxiv.org/abs/1412.6980> (accessed on 6 November 2019).

