



Density-matrix renormalization group: a pedagogical introduction

G. Catarina^{1,2,3,a}  and Bruno Murta^{1,2}

¹ Theory of Quantum Nanostructures Group, International Iberian Nanotechnology Laboratory (INL), 4715-330 Braga, Portugal

² Centro de Física das Universidades do Minho e do Porto, Universidade do Minho, Campus de Gualtar, 4710-057 Braga, Portugal

³ Nanotech@surfaces Laboratory, Empa-Swiss Federal Laboratories for Materials Science and Technology, 8600 Dübendorf, Switzerland

Received 27 April 2023 / Accepted 24 July 2023 / Published online 14 August 2023
© The Author(s) 2023

Abstract. The physical properties of a quantum many-body system can, in principle, be determined by diagonalizing the respective Hamiltonian, but the dimensions of its matrix representation scale exponentially with the number of degrees of freedom. Hence, only small systems that are described through simple models can be tackled via exact diagonalization. To overcome this limitation, numerical methods based on the renormalization group paradigm that restrict the quantum many-body problem to a manageable subspace of the exponentially large full Hilbert space have been put forth. A striking example is the density-matrix renormalization group (DMRG), which has become the reference numerical method to obtain the low-energy properties of one-dimensional quantum systems with short-range interactions. Here, we provide a pedagogical introduction to DMRG, presenting both its original formulation and its modern tensor-network-based version. This colloquium sets itself apart from previous contributions in two ways. First, didactic code implementations are provided to bridge the gap between conceptual and practical understanding. Second, a concise and self-contained introduction to the tensor-network methods employed in the modern version of DMRG is given, thus allowing the reader to effortlessly cross the deep chasm between the two formulations of DMRG without having to explore the broad literature on tensor networks. We expect this pedagogical review to find wide readership among students and researchers who are taking their first steps in numerical simulations via DMRG.

1 Introduction

Understanding the properties of quantum matter is one of the key challenges of the modern era [1]. The difficulties encountered are typically twofold. On the one hand, there is the challenge of modeling all the interactions of a complex quantum system. On the other hand, even when an accurate model is known, solving it is generally not an easy task. In what follows, we will overlook the first challenge and consider only quantum systems for which we can write a model Hamiltonian. Whether such a model is a good description of the physical system or not is thus beyond the scope of this colloquium.

Quantum problems can be divided into two classes: single-body and many-body. In the single-body case, the model Hamiltonian does not include interactions between different quantum particles. In other words, the quantum system can be described as if there was only one quantum particle subject to some potential.

Single-body problems are easy to solve by numerical means, as the dimension of the corresponding Hamiltonian matrix scales linearly with the number of degrees of freedom. For instance, if we consider one electron in N_o spin-degenerate molecular orbitals, we have $2N_o$ possible configurations, as the electron can have either spin- \uparrow or spin- \downarrow in each of the molecular orbitals. In contrast to the single-body case, quantum many-body problems entail interactions between the different quantum particles that compose the system. In that case, the Hamiltonian matrix must take all particles into account, which leads to an exponential growth of its dimension with the number of degrees of freedom. Using the previous example, the basis of the most general many-body Hamiltonian should have 4^{N_o} terms, since every molecular orbital can be empty, doubly-occupied, or occupied by one electron with either spin-up or spin-down. Even if we fix the filling level N_e/N_o (where N_e denotes the number of electrons), we obtain $\binom{2N_o}{N_e}$ con-

^a e-mail: goncalo.catarina@empa.ch (corresponding author)

figurations, which still scales exponentially with N_o .¹ Hence, the exact diagonalization of quantum many-body problems is limited to small systems described by simple models. This is known as the exponential wall problem [2].

To circumvent the exponential wall in quantum mechanics, several numerical methods, each involving a different set of approximations, have been devised. Notable examples are the mean-field approximation, perturbation theory, the configuration interaction method [3], density-functional theory [4–6], quantum Monte Carlo [7], and quantum simulation [8–10], each of which having its own limitations. Additionally, there is the density-matrix renormalization group (DMRG), introduced in 1992 by Steven R. White [11, 12]. This approach, founded on the basis of the variational principle, rapidly established itself as the reference numerical method to obtain the low-energy properties of one-dimensional (1D) quantum systems with short-range interactions [13]. Importantly, a few years after its discovery, DMRG was reformulated in the language of tensor networks [14–16], which allowed for more efficient code implementations [17, 18]. The connection between the original formulation of DMRG and its tensor-network version is by no means straightforward, as the latter involves a variational optimization of a wave function represented by a matrix product state (MPS), making no direct reference to any type of renormalization technique.

The goal of this colloquium is to present a pedagogical introduction to DMRG in both the original and the MPS-based formulations. Our contribution should, therefore, add up to the vast set of DMRG reviews in the literature [13, 16, 19–22]. By following a low-level approach and focusing on learning purposes, we aim to provide a comprehensive introduction for beginners in the field. Bearing in mind that a thorough conceptual knowledge should be accompanied by a notion of practical implementation, we provide as supporting materials simplified and digestible code implementations in the form of documented Jupyter notebooks [23] to put both levels of understanding on firm footing.

The rest of this work is organized as follows. In Sect. 2, we introduce the truncated iterative diagonalization (TID). Although this renormalization technique has been successfully applied to quantum impurity models through Wilson’s numerical renormalization group [24, 25], we illustrate why it is not suitable for the majority of quantum problems. Section 3 contains the original formulation of DMRG, as invented by Steven R. White [11, 12]. We first describe the infinite-system DMRG, which essentially differs from the TID by the type of truncation employed. The truncation

used in DMRG is then shown to be optimal, in the sense that it minimizes the difference between the exact and the truncated wave functions. Importantly, we also clarify the reason that renders this truncation efficient when applied to the low-energy states of 1D quantum systems with short-range interactions. This section ends with the introduction of the finite-system DMRG. In Sect. 4, we give a brief overview on tensor networks, addressing the minimal fundamental concepts that are required to understand how these are used in the context of DMRG. Section 5 shows how, in the framework of tensor networks, the finite-system DMRG can be seen as an optimization routine that, provided a representation of the Hamiltonian in terms of a matrix product operator (MPO), minimizes the energy of a variational MPS wave function. Finally, in Sect. 6, we present our concluding remarks, mentioning relevant topics that are beyond the scope of this review.

In Supplementary Information, we make available a transparent (though not optimized) Python [26] code that, for a given 1D spin model, implements the following algorithms: (i) iterative exact diagonalization, which suffers from the exponential wall problem; (ii) TID; (iii) infinite-system DMRG, within the original formulation. For pedagogical purposes, this code shares the same main structure for the three methods, differing only on a few lines of code that correspond to the implementation of the truncations associated with each method. Following the same didactic approach, we also provide a practical implementation of the finite-system DMRG algorithm in the language of tensor networks.

2 Truncated iterative diagonalization

The roots of DMRG can be traced back to a decimation procedure, to which we refer as TID. Given a large, numerically intractable quantum system, the key idea of this approach is to divide it into smaller blocks that can be solved by exact diagonalization. Combining these smaller blocks together, one at the time and integrating out the high-energy degrees of freedom, this renormalization technique arrives at a description of the full system in terms of a truncated Hamiltonian that can be diagonalized numerically. The underlying assumption of this method is that the low-energy states of the full system can be accurately described by the low-energy states of smaller blocks. The TID routine is one of the main steps in Wilson’s numerical renormalization group [24, 25], which has had notable success in solving quantum impurity problems, such as the Kondo [27] and the Anderson [28] models. As we shall point out below, TID was found to perform poorly for most quantum problems, working only for those where there is an intrinsic energy scale separation, such as quantum impurity models.

We now elaborate on the details of a TID implementation. For that matter, let us consider TID as schematically described in Fig. 1. In the first step, we consider a small system A, with Hamiltonian \mathcal{H}_A , the dimension

¹ For reference, note that for $N_o = N_e = 6$, as we would have in the simplest model for a benzene molecule, there are 924 electronic configurations, which could be encoded in 4 kB of computer memory. However, in the case of a slightly larger molecule such as triangulene, for which $N_o = N_e = 22$, there are roughly 2×10^{12} configurations, which would require 8 TB.

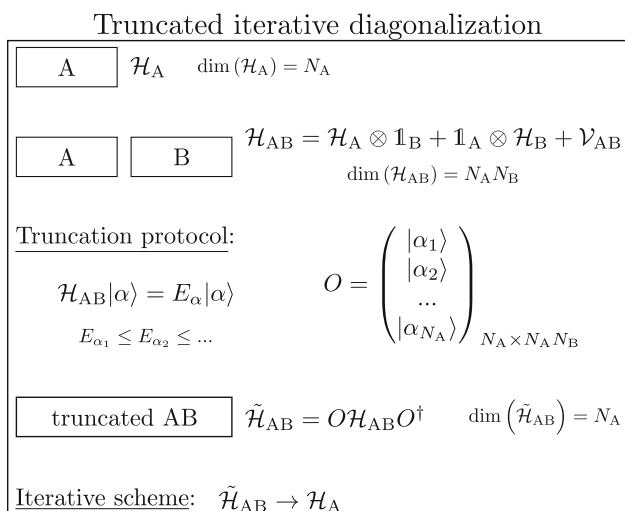


Fig. 1 Schematic description of the truncated iterative diagonalization method. At every iteration, the system size is increased while maintaining the dimension of the Hamiltonian matrix manageable for numerical diagonalization. This is achieved by projecting the basis of the enlarged system onto a truncated basis spanned by its lowest-energy eigenstates. The underlying assumption of this renormalization technique is that the low-energy states of the full system can be accurately described by the low-energy states of smaller blocks

of which, N_A , is assumed to be manageable by numerical means. In the next step, we increase the system size, forming what we denote by system AB, the Hamiltonian of which, \mathcal{H}_{AB} , has dimension $N_A N_B$ and is also assumed to be numerically tractable. The Hamiltonian \mathcal{H}_{AB} includes the Hamiltonians of the two individual blocks A and B, as well as their mutual interactions \mathcal{V}_{AB} . Importantly, if we iterated the procedure at this step, it would be equivalent to doing exact diagonalization, in which case we would rapidly arrive at the situation where the dimension of the Hamiltonian matrix would increase to values that are too large to handle. Instead, in the third step, we diagonalize \mathcal{H}_{AB} and keep only its N_A lowest-energy eigenstates.² These are used to form a rectangular matrix O , which can be employed to project the Hilbert space of the system AB onto a truncated basis spanned by its N_A lowest-energy eigenstates, thereby integrating out the remaining higher energy degrees of freedom. As a consequence, it is possible to find an effective truncated version of any relevant operator defined in the system AB. In particular, we can truncate \mathcal{H}_{AB} , obtaining an effective Hamiltonian $\tilde{\mathcal{H}}_{AB}$ with reduced dimension N_A , which can be used as the input for the first step of the next iteration. This procedure is then iterated until the desired system size is reached. As a final remark, we note that the matrices O should be saved in memory at every iteration, as they

² More generally, the number of kept states does not need to be equal to N_A , but only small enough to stop the exponential growth and make the next iteration manageable.




Fig. 2 Illustration of the failure of the truncated iterative diagonalization method for the problem of a quantum particle in a box. The dashed blue (red) lines represent the two lowest-energy wave functions in the box A (B). The solid black line represents the lowest-energy wave function in the box AB. It is apparent that the lowest-energy state of the larger box cannot be obtained as a linear combination of a few low-energy states of the smaller boxes, thus leading to the breakdown of the principle that underpins the TID approach

are required to obtain the terms \mathcal{V}_{AB} , which we usually only know how to write in the original basis, as well as to compute expectation values of observables.

Despite its rather intuitive formulation, TID turned out to yield poor results for most quantum many-body problems [12]. In fact, White and Noack realized [29] that this renormalization approach could not even be straightforwardly applied to one of the simplest (single-body) problems in quantum mechanics: the particle-in-a-box model (Fig. 2). Even though White and Noack managed to fix this issue by considering various combinations of boundary conditions, this observation was a clear drawback to the aspirations of TID, which motivated the search for a different method. This culminated in the invention of DMRG, which is the focus of the next section.

3 Original formulation of DMRG

3.1 Infinite-system algorithm

In 1992, Steven R. White realized that the eigenstates of the density matrix are more appropriate to describe a quantum system than the eigenstates of its Hamiltonian [11]. This is the working principle of DMRG. In this subsection, we consider the so-called infinite-system DMRG algorithm. Even though it is possible to further improve this implementation scheme,³ it is

³ In this review, we shall focus on finite 1D systems with open boundary conditions. In this regard, the approximate ground state found via infinite-system DMRG by growing the system size at every iteration is refined by finite-system DMRG by sweeping across the fixed-size system, as described in Sect. 3.2. However, for reference, we refer the reader to alternative implementations of DMRG that target the thermodynamic limit, namely, the infinite time-evolving

an instructive starting point as it already contains the core ideas of DMRG. Below, we introduce it in four steps. First, we describe how to apply it, providing no motivation for its structure. Second, we show, on the basis of the variational principle, that the truncation protocol prescribed by this method is optimal. Third, we address its efficiency—i.e., how numerically affordable the truncation required for an accurate description of a large system is—clarifying the models for which it is most suitable. Fourth, we provide a pedagogical code implementation and discuss the results obtained.

3.1.1 Description

The infinite-system DMRG algorithm is schematically described in Fig. 3. In the first step, we consider two blocks, denoted as S (system) and E (environment). As we shall see, both blocks are part of the full system under study, so their designation is arbitrary. Then, we increase the system size by adding two physical sites, one to each block, forming what we denote by blocks L (left) and R (right). We proceed by building the block SB (superblock), which amounts to bundling the blocks L and R. The block SB is the representation of the full system that we intend to describe at every iteration. It should be noted that all block aggregations imply that we account for the individual Hamiltonians of each block, plus their mutual interactions. Finally, we move on to the truncations. As a side remark, we point out that, if we truncated the blocks L and R using the corresponding low-energy states, forming new blocks S and E to use in the first step of the next iteration, this algorithm would be essentially equivalent to TID. Alternatively, we diagonalize the block SB, and use one of its eigenstates $|\psi\rangle$ to build the density matrix $\rho = |\psi\rangle\langle\psi|$.⁴ Then, we compute the reduced density matrices in the subspaces of the blocks L and R, $\sigma_{L/R} = \text{Tr}_{R/L}\rho$, diagonalize them, and keep their eigenvectors with highest eigenvalues. These are used to truncate the blocks L and R, forming new blocks S and E that are taken as inputs of the first step in the next iteration. For clarity, we note that, in the first few iterations, we may skip the truncation protocol (or, equivalently, keep all the eigenvectors of $\sigma_{L/R}$); in that case, the algorithm is equivalent to exact diagonalization, since the block SB is defined in the full Hilbert space. Due to the exponential wall problem, truncations will be required at some iteration; from then on, the Hamiltonian of the block SB is no longer exact, as it is represented in a truncated basis set.

⁴Footnote 3 continued

block decimation [30], a modern version of the infinite-system DMRG [31], and the variational uniform matrix product state algorithm [32].

Here, we should choose the eigenstate $|\psi\rangle$ that we intend to obtain. Most often, it will be the ground state or one of the lowest-energy eigenstates. It is also possible to consider multiple target states $|\psi_n\rangle$, taking $\rho = \sum_n c_n |\psi_n\rangle\langle\psi_n|$, with $\sum_n c_n = 1$. Drawbacks and best practices of this strategy are briefly discussed in Ref. [13].




Fig. 3 Schematic description of the infinite-system DMRG algorithm (see also Fig. 18 of Appendix D for the corresponding pseudocode). Similarly to the TID approach, the system size is increased at every iteration while preventing an exponential growth of the dimension of its Hamiltonian matrix. The truncation employed involves the diagonalization of reduced density matrices, as their eigenvectors with highest eigenvalues are used to obtain an effective description of the enlarged system in a reduced basis. As shown in Sect. 3.1.2, this truncation protocol is optimal and can, in principle, be applied to obtain the best approximation of any state $|\psi\rangle$ of an arbitrary quantum model. In practice, however, this method is mostly useful to probe the low-energy states of 1D quantum problems with short-range interactions (see Sect. 3.1.3)

3.1.2 Argument for truncation

Here, we justify the truncation strategy prescribed above. For that matter, let us consider an exact wave function of the block SB, written as

$$|\psi\rangle = \sum_{i_L=1}^{N_L} \sum_{i_R=1}^{N_R} \psi_{i_L, i_R} |i_L\rangle \otimes |i_R\rangle, \quad (1)$$

where $|i_L\rangle$ ($|i_R\rangle$) denotes a complete basis of the block L (R), with dimension N_L (N_R). We now propose a

variational wave function of the form

$$|\tilde{\psi}\rangle = \sum_{\alpha_L=1}^{D_L} \sum_{i_R=1}^{N_R} c_{\alpha_L, i_R} |\alpha_L\rangle \otimes |i_R\rangle, \tag{2}$$

where $|\alpha_L\rangle$ denotes a truncated basis of the block L, with reduced dimension $D_L < N_L$. The goal is to find the states $|\alpha_L\rangle$ and the variational coefficients c_{α_L, i_R} that provide the best approximation of the truncated wave function $|\tilde{\psi}\rangle$ to the exact wave function $|\psi\rangle$, for a given D_L . This can be achieved by minimizing $\| |\psi\rangle - |\tilde{\psi}\rangle \|^2$.

The exact wave function is normalized, i.e., $\langle \psi | \psi \rangle = 1$. Using this property, we obtain

$$\begin{aligned} \| |\psi\rangle - |\tilde{\psi}\rangle \|^2 = 1 - \sum_{i_L, i_R, \alpha_L} \left(\psi_{i_L, i_R}^* c_{\alpha_L, i_R} \langle i_L | \alpha_L \rangle \right. \\ \left. + c_{\alpha_L, i_R}^* \psi_{i_L, i_R} \langle \alpha_L | i_L \rangle \right) + \sum_{\alpha_L, i_R} |c_{\alpha_L, i_R}|^2, \end{aligned} \tag{3}$$

where we have also used the orthonormal properties of the basis states, e.g., $\langle i_R | i'_R \rangle = \delta_{i_R, i'_R}$. To minimize the previous expression, we impose that its derivative with respect to the variational coefficients c_{α_L, i_R} (or c_{α_L, i_R}^*) must be zero. This leads to

$$c_{\alpha_L, i_R} = \sum_{i_L=1}^{N_L} \psi_{i_L, i_R} \langle \alpha_L | i_L \rangle. \tag{4}$$

Inserting Eqs. (4) into (3), we obtain

$$\| |\psi\rangle - |\tilde{\psi}\rangle \|^2 = 1 - \sum_{\alpha_L=1}^{D_L} \langle \alpha_L | \sigma_L | \alpha_L \rangle, \tag{5}$$

where we have introduced the reduced density matrix of the state $|\psi\rangle$ in the subspace of block L

$$\sigma_L = \text{Tr}_R \rho = \sum_{i_R=1}^{N_R} \langle i_R | \rho | i_R \rangle, \tag{6}$$

defined in terms of the full density matrix

$$\rho = |\psi\rangle \langle \psi|. \tag{7}$$

Looking at Eq. (5), we observe that it involves a partial trace of the reduced density matrix σ_L (note that σ_L is an $N_L \times N_L$ matrix, but the sum over α_L runs over D_L terms only). Since σ_L is a density matrix, its full trace must be equal to 1, in which case the minimization of $\| |\psi\rangle - |\tilde{\psi}\rangle \|^2$ is accomplished by maximizing the partial trace of σ_L . Per the Schur–Horn theorem [33,34], the states $|\alpha_L\rangle$ that accomplish this are those that diagonalize σ_L with highest eigenvalues λ_{α_L} (which are all

non-negative, since any density matrix is positive semi-definite), that is

$$\sigma_L |\alpha_L\rangle = \lambda_{\alpha_L} |\alpha_L\rangle, \quad \lambda_1 \geq \lambda_2 \geq \dots, \tag{8}$$

thus leading to

$$\| |\psi\rangle - |\tilde{\psi}\rangle \|^2 = 1 - \sum_{\alpha_L=1}^{D_L} \lambda_{\alpha_L}. \tag{9}$$

Let us now put into words what we have just demonstrated. Starting from an exact wave function $|\psi\rangle$, we can obtain a truncated (in the subspace of the block L) wave function $|\tilde{\psi}\rangle$ that best approximates $|\psi\rangle$ by going through the following protocol. First, we build the density matrix $\rho = |\psi\rangle \langle \psi|$ and compute the reduced density matrix $\sigma_L = \text{Tr}_R \rho$. Then, we diagonalize σ_L and form a $D_L \times N_L$ matrix O whose lines are the eigenvectors of σ_L with highest eigenvalue. Finally, $|\tilde{\psi}\rangle$ is obtained as $|\tilde{\psi}\rangle = O|\psi\rangle$. Repeating the same strategy for the block R, for which the derivation is completely analogous, we arrive at the truncation scheme described in Sect. 3.1.1.

The calculation of $\| |\psi\rangle - |\tilde{\psi}\rangle \|^2$ at every iteration of the algorithm, using Eq. (9), can be used as a measure of the quality of the corresponding truncation. Therefore, instead of fixing a given D_L , we can impose a maximum tolerance for $\| |\psi\rangle - |\tilde{\psi}\rangle \|^2$, obtaining an adaptive truncation scheme. As a final remark, we note that, while the general derivation presented here applies to any state $|\psi\rangle$ of an arbitrary quantum problem, the efficiency of DMRG relies on how large D_L must be to ensure that the truncation does not compromise the accurate quantitative description of the system under study. This subject is addressed below.

3.1.3 Efficiency

Recalling Eq. (9), it is apparent that the efficiency of DMRG relies on how fast the eigenvalues of the reduced density matrices decay for the quantum state $|\psi\rangle$ under study. However, this property is, in general, unknown.⁵ Instead, the entanglement entropy—for which general results are known or conjectured [37]—can be used as a proxy, as explained below.

The blocks L and R form a bipartition of the full system, represented by the block SB. We can therefore define the von Neumann entanglement entropy (of the state $|\psi\rangle$) between L and R as

$$\begin{aligned} \mathcal{S} &\equiv \mathcal{S}(\sigma_L) = -\text{Tr}(\sigma_L \log_2 \sigma_L) \\ &= \mathcal{S}(\sigma_R) = -\text{Tr}(\sigma_R \log_2 \sigma_R). \end{aligned} \tag{10}$$

⁵ For completeness, we note that, for 1D systems at quantum critical points, the distribution of the eigenvalues of the reduced density matrix can be obtained from a conformal field theory [35], but it is only strictly valid in the continuum limit. Nevertheless, this result is often referenced in the study of the complexity of the MPS representation of 1 + 1D critical ground states (see, e.g., Ref. [36]).




Fig. 4 Relation between dimensionality and range of interactions on a lattice model. In the example depicted, a 3×3 two-dimensional square lattice with nearest-neighbor hopping terms is described as a 1D chain with hoppings up to fifth neighbors. In general, the same mapping applied to an $N \times N$ lattice leads to (nonlocal) hopping terms between sites separated by up to $2N - 1$ units of the 1D chain

Focusing on the block L , without loss of generality, we write

$$S = - \sum_{\alpha_L=1}^{N_L} \lambda_{\alpha_L} \log_2 \lambda_{\alpha_L} \simeq - \sum_{\alpha_L=1}^{D_L} \lambda_{\alpha_L} \log_2 \lambda_{\alpha_L}, \tag{11}$$

where we have restricted the sum over α_L to the D_L highest eigenvalues of σ_L . This approximation is valid, since we are fixing D_L , so that $\| |\psi\rangle - |\tilde{\psi}\rangle \|^2 \simeq 0$, which implies, by virtue of Eq. (9), that the remaining eigenvalues are close to zero; given that $\lim_{\lambda \rightarrow 0^+} \lambda \log_2 \lambda = 0$, it follows that the lowest eigenvalues of σ_L can be safely discarded in the calculation of the entanglement entropy. Within this assumption, it is also straightforward to check that S is maximal if $\lambda_{\alpha_L} = 1/D_L$, $\alpha_L = 1, 2, \dots, D_L$, which allows us to write

$$S \leq \log_2 D_L, \tag{12}$$

leading to

$$D_L \geq 2^S. \tag{13}$$

Using Eq. (13), we can make a rough estimate of the order of magnitude of D_L

$$D_L \sim 2^S. \tag{14}$$

The scaling of S with the size of a translationally invariant quantum system is a property that is widely studied. In particular, there are exceptional quantum states that obey the so-called area laws [37], meaning that S , instead of being an extensive quantity,⁶ is at most proportional to the boundary of the two partitions. The area laws are commonly found to hold for the ground states of gapped Hamiltonians with local interactions [37]; this result has been rigorously demonstrated in the 1D case [38]. It should also be noted that, for the ground states of 1D critical/gapless local models, the scenario is not dramatically worse as S is typically verified to scale only logarithmically with the chain length [39, 40].

⁶ Note that this situation, expected in the most general case, leads to an exponential scaling of D_L with the system size, which is impractical for numerical purposes.

In summary, considering the ground state of a local Hamiltonian describing a \mathcal{D} -dimensional system of size \mathcal{L} in each dimension, we expect to have:

- $S \sim \text{const.}$, for 1D gapped systems. This implies a favorable scaling $D_L \sim 2^{\text{const.}}$.
- $S \sim c \log_2 \mathcal{L}$, for 1D gapless models. This leads to $D_L \sim 2^{c \log_2 \mathcal{L}}$, yielding a power law in \mathcal{L} , which is usually numerically manageable in practical cases.
- $S \sim \mathcal{L}^{\mathcal{D}-1}$, for gapped systems in $\mathcal{D} = 2, 3$ dimensions. This implies $D_L \sim 2^{\mathcal{L}^{\mathcal{D}-1}}$, resulting in an exponential scaling that severely restricts the scalability of numerical calculations.

In short, we see that the truncation strategy employed in DMRG is in principle suitable for 1D quantum models (gapped or gapless), but not in higher dimensions. Notable exceptions are two-dimensional problems whose solutions can be obtained or extrapolated from lattices where the size along one of the two dimensions is rather small, such as stripes or cylinders (see Ref. [41] for a review on the use of DMRG to study two-dimensional systems). In fact, there is a relation between dimensionality and range of interactions in finite systems (Fig. 4), from which it also becomes apparent that DMRG is in practice only efficient when applied to models with short-range interactions. Finally, it is reasonable to expect that the previous statements may hold not only for ground states but also for a few low-lying states.

3.1.4 Code implementation

In Supplementary Information, we present a didactic code implementation of the infinite-system DMRG algorithm, also made available at https://github.com/GCatarina/DMRG_didactic. In this documented Jupyter notebook, written in Python, we focus on tackling spin-1 Heisenberg chains with open boundary conditions. The generalization to different spin models is completely straightforward. As for other types of quantum problems (e.g., fermionic models), this code can be readily used after simply defining the operators that appear in the corresponding Hamiltonian. We also note that a slight modification of the algorithm has been proposed to better deal with periodic boundary conditions [43].

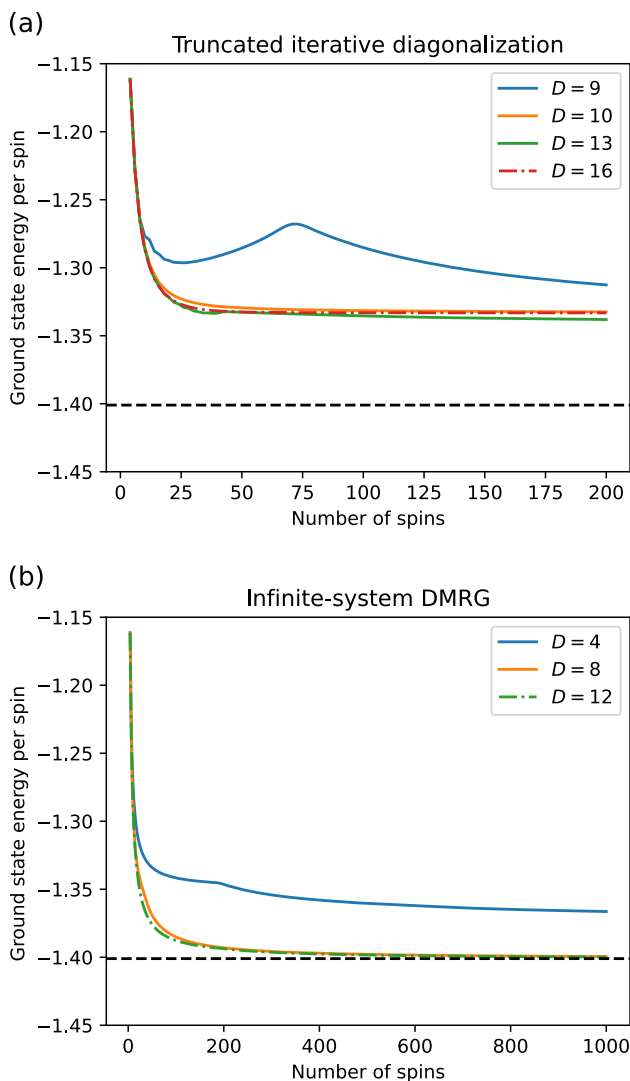


Fig. 5 Benchmark results of truncated iterative diagonalization and infinite-system DMRG methods applied to open-ended spin-1 Heisenberg chains. Ground-state energy per spin, as a function of the number of spins, obtained with TID (a) and infinite-system DMRG (b), for different values of D , which reflects the truncation employed, as described in the text. In both algorithms, every iteration implies the diagonalization of a Hamiltonian matrix of maximal dimension $9D^2 \times 9D^2$. Larger matrices are allowed if degeneracies to within numerical precision are found at the truncation threshold, as explained in the code documentation. The dashed black line marks the known result in the thermodynamic limit [42]

For pedagogical purposes, our Jupyter notebook is structured in three parts. First, we adopt the scheme described in Fig. 3, but make no truncations. This is the same as doing exact diagonalization. It is observed that, at every iteration, the running time of the code increases dramatically, reflecting the exponential wall problem. Second, maintaining the same scheme, we make a truncation where the D lowest-energy states of the block L (R) are used to obtain the new block S




Fig. 6 Magnetic properties of spin-1 Heisenberg chains computed by infinite-system DMRG. Local distribution of magnetization for the ground state with quantum number $S_z = +1$ (where S_z denotes the total spin projection) of an open-ended chain composed of 100 spins. The calculated local moments are exponentially localized at both edges of the chain, reflecting the fractionalization of the ground state into two effective spin-1/2 edge states. These results were obtained with an adaptive implementation in which the truncation error at every iteration was imposed to be below 10^{-4} . A small Zeeman term was added to the Hamiltonian to target the $S_z = +1$ ground state

(E). This is equivalent to the TID approach. In Fig. 5a, we plot the ground state energy per spin, as a function of the number of spins, obtained with this strategy, for different values of D . Our calculations show a disagreement of at least 5% with the reference value [42], which does not appear to be overcome by considering larger values for D . Therefore, we conclude that TID is not fully reliable for this problem. Third, we implement the infinite-system DMRG, where we first set a fixed value for $D \equiv D_L = D_R$ in the truncations. Computing the ground state energy per spin with this method, the results obtained are very close to the reference value, even for small values of D , as shown in Fig. 5b. For completeness, we also implement an adaptive version of the algorithm where the values of $D_{L/R}$ used at every iteration are set as to keep the truncation error, given by Eq. (9), below a certain threshold. This adaptive implementation is used to compute the expectation values presented in Fig. 6, which show a known signature of the emergence of fractional spin-1/2 edge states in the model [42].

3.2 Finite-system scheme

Within the infinite-system DMRG approach, the size of the system that we aim to describe increases at every iteration of the algorithm. Therefore, the wave function targeted at each step is differential. This can lead to a poor convergence of the variational problem or even to incorrect results. For instance, a metastable state can be favored by edge effects in the early DMRG steps, where the embedding with the environment is not so

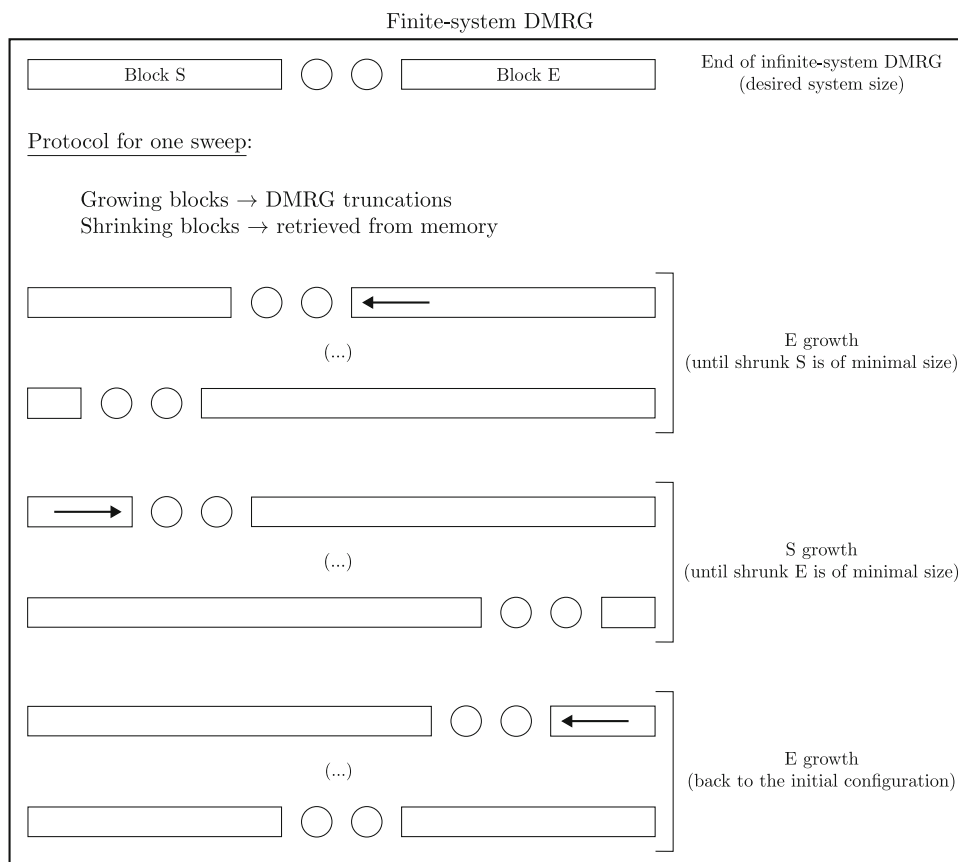


Fig. 7 Breakdown of the finite-system DMRG routine. At the first stage, infinite-system DMRG is used to obtain an effective description for the target wave function of a system with desired size. This is followed by a sweeping protocol where one of the blocks is allowed to grow while the other is shrunk, thus keeping the total system size fixed. To prevent the exponential scaling, DMRG truncations (targeting the intended state) are employed for the growing blocks. The shrinking blocks are retrieved from memory, using stored data of the latest description of the block with such size (either from the infinite-system routine or from an earlier step of the sweeping procedure). The growth direction is reversed when the shrinking block reaches its minimal size. A typical strategy is to fix a maximal truncation error for the DMRG truncations, and perform sweeps until convergence in energy (and/or other physical quantities of interest) is attained; this approach ensures that the description of the target wave function is improved (or at least not worsened) at each step of the sweeping protocol

effective due to its small size, and the lack of “thermalization” in the following iterations may not allow for a proper convergence to the target state.

In this subsection, we present the so-called finite-system DMRG method, which manages to fix the aforementioned issues to a large extent. The breakdown of this algorithm is shown in Fig. 7. Its first step consists in applying the infinite-system routine to obtain an effective description for the target wave function of a system with desired size. Then, a sweeping protocol is carried out to improve this description. In this part, one of the blocks is allowed to grow, while simultaneously shrinking the other, thus keeping the overall system size fixed. DMRG truncations (targeting the intended state) are employed for the growing blocks, whereas the shrinking blocks are retrieved from previous steps. When the shrinking block reaches its minimal size, the growth direction is reversed. A complete loop of this protocol, referred to as a *sweep*, entails the shrinkage of the two blocks to their minimal sizes, and the return to

the initial block configuration. For a fixed truncation error, every step of a sweep must lead to a better (or at least equivalent) description of the target wave function; when the target is the ground state, this implies a variational optimization in which the estimated energy is a monotonically non-increasing function of the number of sweep steps performed. This property is at the heart of the MPS formulation of DMRG (see Sect. 5.1).

As a final remark, we wish to clarify a few subtleties related to the variational character of DMRG. For that matter, let us focus on the case where the target is the ground-state wave function. According to the derivation presented in Sect. 3.1.2, it is straightforward to check that the DMRG truncations are variational in the number of kept states: a larger value of $D_{L/R}$ implies a better (or at least equivalent) description of the exact wave function, and, hence, a non-increasing energy estimation. On top of that, we have just argued that, as long as we keep a fixed truncation error, the finite-system method is also variational in the number of

sweeps. Hence, the finite-system algorithm has an additional knob of optimization—the number of sweeps—that allows to improve the results of the infinite-system scheme.

4 Tensor-network basics

The modern formulation of DMRG is built upon tensor networks [14–16]. Indeed, virtually all state-of-the-art implementations of DMRG [17, 18] make use of MPSs and MPOs. Although pedagogical reviews on these and other tensor networks are available [44–46], their scope goes far beyond DMRG, as they provide the reader with the required background to explore the broader literature on tensor-network methods. Here, we take a more focused approach, giving the minimum necessary framework on tensor networks to understand the MPS-based version of the finite-system DMRG algorithm, which is discussed in detail in Sect. 5.

4.1 Diagrams and key operations

A tensor can be simply regarded as a mathematical object that stores information in a way that is determined by the number of indices $r \in \mathbb{N}^0$ (referred to as the *rank* of the tensor), their dimensions $\{d_i\}_{i=1}^r$ (i.e., the i^{th} index can take $d_i \in \mathbb{N}^+$ different values), and the order by which those indices are organized. The total number of entries of a tensor is $\prod_{i=1}^r d_i$. The most familiar examples of tensors are scalars (i.e., rank-0 tensors, each corresponding to a single number, thus not requiring any labels), vectors (i.e., rank-1 tensors, where every value is labeled by a single index that takes as many different values as the size of the vector), and matrices (i.e., rank-2 tensors, where every entry is characterized by two indices, one labeling the rows and another the columns). In general, each number stored in a rank- r tensor is labeled in terms of an ordered array of r indices, which can be regarded as its coordinates within the structure of the tensor. In Fig. 8a–c, we show how tensors are represented diagrammatically.

Although the number of indices, their dimensions, and the order by which they are organized are crucial to unambiguously label the entries of a tensor, these properties—to which we shall refer as the *shape* of the tensor—are immaterial in the sense that we can fuse, split, or permute its indices without actually changing the information contained within it. For clarity, let us consider the following 2×4 matrix $A_{\alpha\beta}$ with $\alpha \in \{0, 1\}$, $\beta \in \{0, 1, 2, 3\}$:

$$\alpha \text{---} \bigcirc \text{---} \beta = \begin{pmatrix} 5 & 0 & 7 & 8 \\ 4 & 9 & 3 & 2 \end{pmatrix}.$$

We can *reshape* this rank-2 tensor by *fusing* its two indices, yielding the 8-dimensional vector $A_{(\alpha,\beta)} \equiv A_\gamma$ if the row index α is chosen to precede the column index β

$$(\alpha, \beta) \equiv \gamma \text{---} \bigcirc = (5 \ 0 \ 7 \ 8 \ 4 \ 9 \ 3 \ 2),$$

or $A_{(\beta,\alpha)} \equiv A_\delta$ if β takes precedence

$$(\beta, \alpha) \equiv \delta \text{---} \bigcirc = (5 \ 4 \ 0 \ 9 \ 7 \ 3 \ 8 \ 2).$$

Likewise, we can *split* the $d = 4$ column index β into two $d = 2$ indices, $\beta_1, \beta_2 \in \{0, 1\}$, corresponding to the least and most significant digits of the binary decomposition of β , respectively. This yields the rank-3 tensor

$$\alpha \text{---} \bigcirc \text{---} \beta_1 \text{---} \beta_2 = \begin{pmatrix} 5 & 0 \\ 4 & 9 \end{pmatrix}, \quad \alpha \text{---} \bigcirc \text{---} \beta_1 \text{---} \beta_2 = \begin{pmatrix} 7 & 8 \\ 3 & 2 \end{pmatrix},$$

where β_2 corresponds to the rightmost leg in the diagrams above. Alternatively, we can leave the rank unchanged, *permuting* the row and column indices to yield the 4×2 transpose matrix

$$\beta \text{---} \bigcirc \text{---} \alpha = \begin{pmatrix} 5 & 4 \\ 0 & 9 \\ 7 & 3 \\ 8 & 2 \end{pmatrix}.$$

In all three cases, even though we end up with tensors of different shape, all of them store exactly the same content as the original matrix, albeit in a different way. This is the key point: reshaping a tensor (by fusing or splitting indices) or simply permuting its indices merely restructures how the information is stored, leaving the information itself unaffected. In the context of numerical implementations, we note that these tensor operations can be applied to arbitrary-rank tensors via standard built-in functions (e.g., `numpy.reshape` and




Fig. 8 Diagrammatic representation of simple examples of tensors: **a** vector (i.e., rank-1 tensor), **b** matrix (i.e., rank-2 tensor), and **c** rank-3 tensor. Tensor networks are constructed by joining individual tensors, which is accomplished by contracting (i.e., summing over) indices in common. **d** Example of contraction between rank-2 and rank-3 tensors. Common index j is contracted. Free indices i , k , and l are represented through open legs. **e** Example of canonical tensor network, MPS. Each local tensor has one free index. There is one contracted index (also known as bond) between every pair of adjacent tensors. **f** Representation of generic rank- N tensor

`numpy.transpose` in Python). The time complexity of reshaping a tensor or permuting its indices is essentially negligible, as these operations just modify a flag associated with the tensor that defines its shape rather than actually moving its elements around in memory.

Thus far, we have only considered isolated tensors. However, based on the diagrammatic representations illustrated in Fig. 8a–c, where each index corresponds to a leg, we can think of joining two individual tensors by linking a pair of legs, one from each tensor, as shown in Fig. 8d. Algebraically, such link/bond corresponds to a sum over a common index shared by the two tensors; the outcome of this operation can be explicitly obtained in Python via `numpy.einsum`. Of course, this process can be generalized to an arbitrary number of tensors, resulting in *tensor networks* of arbitrary shapes and sizes. Here, we will focus on the so-called *matrix product states* (MPSs), relevant for DMRG. A diagram of an MPS is shown in Fig. 8e; it comprises both free indices (i.e., open legs) and contracted indices (i.e., bonds). The elements of an MPS are uniquely identified by the free indices, but, unlike the case of an isolated tensor, their values are not immediately available, as the contracted indices must be summed over to obtain them. In the context of DMRG, an MPS with N free/physical indices is typically used to represent a quantum state of a system with N sites.

Even though the order by which sums over contracted indices are performed does not affect the obtained result, different orders may produce substantially different times of execution, especially if the tensor networks in question are large. For the 1D tensor networks herein considered, the type of contractions that we need to deal with are essentially those shown in Fig. 9, for which there are two possible contraction strategies. Contracting multiple bonds of a tensor network essentially amounts to performing nested loops. When we sum over a given contracted index, corresponding to

the current innermost loop, we effectively have to fix the dummy variables of the outer loops. However, all possible values that such dummy variables can take must be considered. In the scheme of Fig. 9a, we first contract the D -dimensional bond linking tensors B and C , which involves order $\mathcal{O}(D)$ operations on its own, but we must repeat this for all possible combinations of values of all other indices of tensors B and C , which are $\mathcal{O}(D^4)$, yielding a total scaling of $\mathcal{O}(D^5)$. The second step contracts both bonds linking A to BC , taking $\mathcal{O}(D^4)$ operations. For Fig. 9b, in turn, contracting first the bond between A and B takes $\mathcal{O}(D^4)$ operations, and the same scaling is obtained for the second step. Hence, (b) has an overall cost of $\mathcal{O}(D^4)$, which is more favorable than the $\mathcal{O}(D^5)$ scaling of (a). In general, the problem of determining the optimal contraction scheme is known to be NP-hard [47, 48], but this issue only arises in two and higher dimensions. For our purposes, the cases described above are all we need to know about tensor-network contractions.

Tensor networks can be regarded as tensors with internal structure. Therein lies their great virtue: such internal structure allows for a compact storage of information, which greatly reduces the memory requirements of the variational methods that use these tensor networks as their ansätze. For concreteness, let us compare the N -site MPS shown in Fig. 8e to an isolated rank- N tensor (resultant, e.g., from contracting all the bonds of the N -site MPS), as shown in Fig. 8f. Assuming that free and contracted indices have dimension d and D , respectively, while the isolated tensor requires storing a total of d^N numbers in memory, the MPS only involves saving the entries of $N-2$ rank-3 $D \times d \times D$ tensors in the bulk and 2 rank-2 $d \times D$ tensors at the ends, yielding $\mathcal{O}(ND^2d)$ numbers saved in memory. In other words, the memory requirements of methods based on MPSs scale linearly with the system size N , in con-





Fig. 9 Comparison of two strategies to contract a tensor network comprising three tensors. All indices, both free and contracted, are assumed to have dimension D for the purpose of estimating scaling of cost of contractions. **a** First, index γ linking tensors B and C is contracted, and then, indices α and β are summed over, yielding an overall cost of $\mathcal{O}(D^5)$. **b** First, index α linking tensors A and B is contracted, and then, indices β and γ are summed over, resulting in $\mathcal{O}(D^4)$ cost. Even though both strategies yield the same outcome, **b** is preferred, since its execution time scales more favorably with the index dimension D

trast with the exponential scaling associated with an unstructured tensor.


4.2 Singular value decomposition

The success of the original formulation of DMRG in tackling quantum many-body problems in a scalable way rests upon the projection of the Hilbert space onto the subspace spanned by the highest-weight eigenstates of the reduced density matrix on either side of the bipartition considered. In the MPS-based formulation, the analog operation (see Sect. 5.2) corresponds to the *singular value decomposition* (SVD) of the local tensors that compose the MPS.

SVD consists of factorizing any $m \times n$ real or complex matrix M in the form $M = USV^\dagger$, where U and V are $m \times m$ and $n \times n$ unitary matrices, respectively, and S is an $m \times n$ matrix with non-negative real numbers (some of which possibly zero) along the diagonal and all remaining entries equal to zero



In the schematic representations of SVD above, the parallel horizontal and vertical lines forming the grids within U and V^\dagger serve to illustrate that the respective rows and columns form an orthonormal set, which is the defining property of a unitary matrix. As highlighted by the shaded regions, all entries of the the last $n - m$ columns (if $m < n$) or the last $m - n$ rows (if $m > n$) of S are zero, so we can remove such redundant information by truncating U , S and V^\dagger (the truncated versions of which we write as U , S and V^\dagger) accordingly



This is the so-called *thin* or *reduced* SVD, as opposed to the *full* SVD described earlier. Both are implemented in Python via `numpy.linalg.svd`, setting the Boolean input parameter `full_matrices` appropriately. Henceforth, unless stated otherwise, we shall consider the thin SVD, as it yields the most compact factorization of the original matrix M .

A brief overview of some terminology related to SVD is in order. First, in the thin SVD diagrams above, V^\dagger in the $m < n$ case and U in the $m > n$ case are rectangular

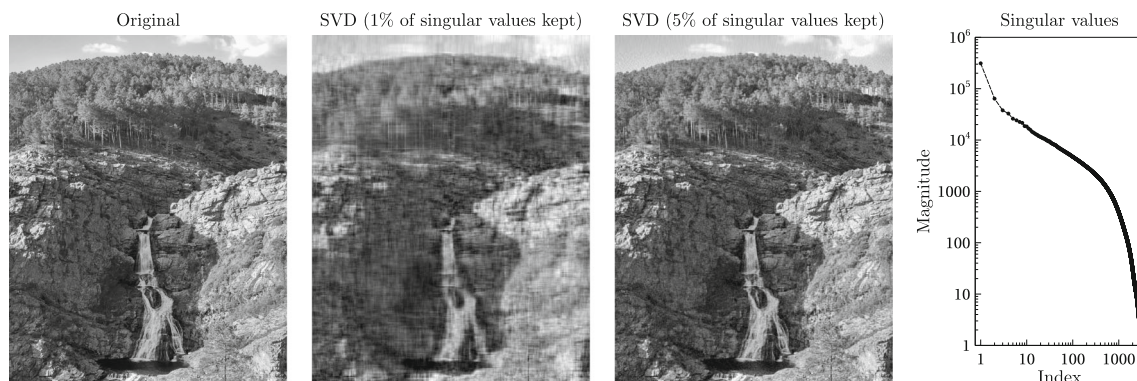


Fig. 10 Singular value decomposition for image compression. The original photo (taken at Físgas de Ermelo, Portugal) is stored as a 3335×2668 matrix, where each entry corresponds to a pixel and the values encode the grayscale color. The compressed images are obtained by applying SVD to this matrix, keeping only the highest singular values (namely, 1% and 5% of the total 2668 singular values). The distribution of the singular values is shown in the rightmost panel

matrices, and therefore, neither is unitary. Nevertheless, as illustrated through the parallel lines, the rows of V^\dagger in the $m < n$ case and the columns of U in the $m > n$ case still form an orthonormal basis, so the former is said to be *right-normalized* (i.e., $V^\dagger(V^\dagger)^\dagger = V^\dagger V = \mathbb{1}$) and the latter *left-normalized* (i.e., $U^\dagger U = \mathbb{1}$). The columns of U and the rows of V^\dagger are referred to as left- and right-singular vectors. The diagonal entries of the $\min\{m, n\} \times \min\{m, n\}$ matrix S are called *singular values*. The Schmidt rank $r_S \leq \min\{m, n\}$ is the number of nonzero singular values. By exploiting the gauge freedom of SVD (see Appendix A), the singular values are conventionally stored in descending order, which is useful when truncations are considered, as explained below.

The application of the thin SVD to a rectangular matrix allows for a trivial truncation of the bond dimensions between the factorized matrices. Further truncations can be implemented by discarding singular values of negligible magnitude. If the discarded singular values are zero, this procedure is exact. Otherwise, some information is lost, but the strategy of discarding the lowest singular values is known to yield the optimal truncation [49, 50]. Therefore, SVD is widely used for data compression, being particularly efficient in cases where the singular values decay rapidly. In Fig. 10, we show such an example, where SVD is used to compress a black-and-white photograph. We observe that, by keeping only the 1% highest singular values, the image obtained already exhibits most of the features of the original photo, though noticeably blurred out. This blur is significantly reduced when the number of kept singular values is increased to just 5%.

A wave function can always be exactly represented by an MPS, although this will generally entail an exponential growth of the bond dimensions from the ends toward the center of the MPS (see Sect. 4.3.3). Within the context of MPS-based DMRG, SVD is adopted both to truncate the bond dimensions of the MPS and to transform it into convenient canonical forms, which we shall introduce in Sect. 4.3.2. However, since SVD is

a linear algebraic method, it applies to matrices and not to the rank-3 tensors found in the non-terminal sites of an MPS. As a result, these tensors have to be reshaped by fusing two indices. There are two possibilities for this, depending on which leg we choose to fuse the physical index with (Fig. 11). In Fig. 11a, we end up with a left-normalized tensor U at the current site, with the remaining SV^\dagger being contracted with the next local tensor to the right of the MPS. In Fig. 11b, the right-normalized tensor V^\dagger is the final form of the tensor at the current site, and US is contracted with the next local tensor to the left. The expressive power of an MPS is determined by the bond dimension cutoff D , which sets the maximum size of the contracted indices (e.g., α , γ , and ϵ in Fig. 11). The dimension d of the physical indices (e.g., β in Fig. 11) is fixed by the local degrees of freedom of the problem under consideration (e.g., $d = 2s + 1$ for a spin- s quantum model). As a result, in Fig. 11a, b, both dimensions of the matrix resulting from reshaping the rank-3 tensor are $\mathcal{O}(D)$. Computing the SVD of an $m \times n$ matrix (with $m > n$) takes $\mathcal{O}(m^2n + n^3)$ floating-point operations [51]. Hence, within the context of MPS-based DMRG, the time complexity of SVD is $\mathcal{O}(D^3)$.

4.3 Matrix product states

This subsection introduces the key operations required to manipulate MPSs. In particular, we discuss how to compute overlaps between two MPSs and expectation values of MPSs for local operators.⁷ Three MPS canonical forms that simplify some of these computations are introduced; the construction of all of them merely involves a sequential sitewise application of SVD, as described in Fig. 11. For completeness, we also explain how to obtain an MPS representation of a general wave function, even though this procedure is not essential for

⁷ A more general discussion of the computation of expectation values with MPSs is deferred to the next subsection, where we introduce MPOs.




Fig. 11 Singular value decomposition of rank-3 tensor A belonging to a matrix product state. **a** Central/physical index β is fused with leftmost index α to yield left-normalized tensor U at current site after SVD and index splitting. The remaining SV^\dagger is contracted with the local tensor that appears to the right of A in the MPS. **b** Central/physical index β is fused with rightmost index γ to yield right-normalized tensor V^\dagger at current site after SVD and index splitting. The remaining US is contracted with the local tensor that appears to the left of A in the MPS. Triangular shapes indicate left- and right-normalization of U and V^\dagger , respectively. Diamond-shaped diagram illustrates that S is diagonal

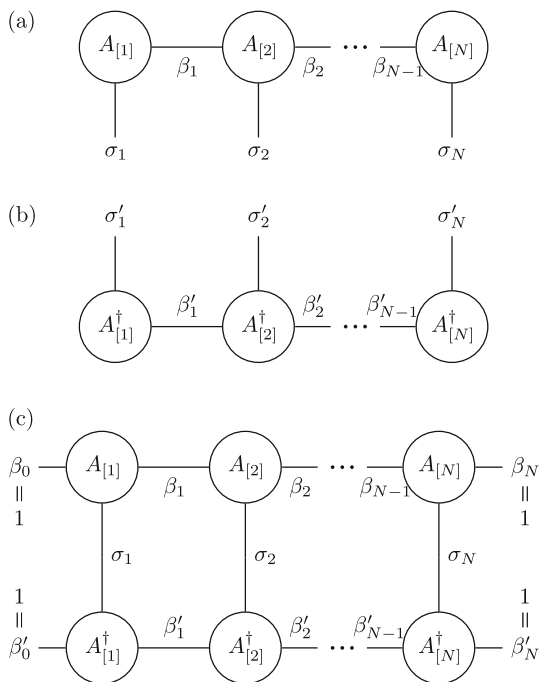


Fig. 12 Diagrammatic representation of **a** MPS for ket $|\psi\rangle$, **b** MPS for bra $\langle\psi|$, and **c** contraction of two previous MPSs to compute norm $\langle\psi|\psi\rangle$. In **c**, singleton dummy indices $\beta_0, \beta'_0, \beta_N$, and β'_N were added on either side of both MPSs to ease discussion of efficient method to contract tensors down to scalar $\langle\psi|\psi\rangle$ (see Fig. 13)


DMRG. In general, we shall consider N -site MPSs with bond dimension D and physical index dimension d .

4.3.1 Overlaps

Using Dirac’s Bra–Ket notation, the MPS representations of a ket $|\psi\rangle$ and its bra $\langle\psi|$ are shown in Fig. 12a, b, respectively. The diagrammatic representation of the norm of this state, $\langle\psi|\psi\rangle$, amounts to linking the two MPSs by joining the physical indices $\{\sigma_i\}_{i=1}^N$, as shown in Fig. 12c. The question, then, is how to contract such

tensor network to arrive at the scalar $\langle\psi|\psi\rangle$. A naïf approach would be to fix the same set of physical indices in the bra and the ket ($\sigma_i = \sigma'_i$), contract the remaining bonds ($N - 1$ at the ket and $N - 1$ at the bra), multiply the scalars obtained in the bra and the ket, and then sum over all possible values of the physical indices. The problem, however, is that $\{\sigma_i\}_{i=1}^N$ take d^N different values, so this would be exponentially costly in N . Fortunately, there is a contraction scheme linear in N that resembles the process of *closing a zipper* [52].

In Fig. 13, we illustrate this closing-the-zipper contraction scheme of the overlap between two MPSs.⁸ The contraction is divided in N steps; at the n^{th} step, the local tensors $A_{[n]}$ and $A_{[n]}^\dagger$ are contracted with the tensor $C_{[n-1]}$ that stores the outcome of all contractions from previous steps, yielding the tensor $C_{[n]}$ to be used in the next step



To make sense of the first and final steps, it is helpful to add singleton dummy indices at each end of the two MPSs, as illustrated in Fig. 12c. This allows to apply the first step of the recursive process depicted in Fig. 13 with $C_{[0]}$ initialized as the 1×1 identity matrix (i.e., the scalar 1). At the N^{th} and final step, the recursive relation results in the rank-2 tensor $C_{[N]}$, with both of its indices β_N and β'_N having trivial dimension 1. This scalar corresponds precisely to the norm $\langle\psi|\psi\rangle$ we were after. Of course, we can cover the tensor network from right to left instead, producing exactly the same outcome. At each step, we make use of the

⁸ For simplicity, we consider the computation of the norm, in which case the bra and ket correspond to the same state. The generalization to the case of an overlap $\langle\phi|\psi\rangle$ between two states $|\phi\rangle$ and $|\psi\rangle$ is straightforward.

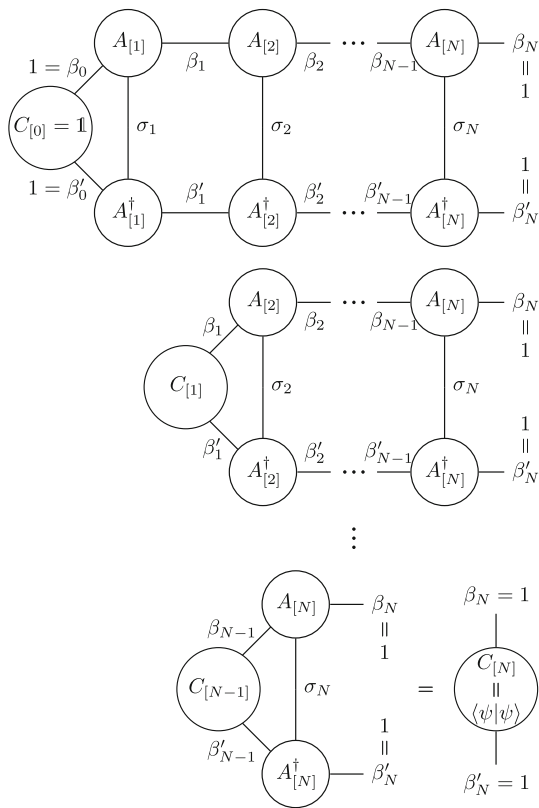


Fig. 13 Schematic description of closing-the-zipper strategy to perform contraction of tensor network resulting from the overlap between two MPSs representing the ket $|\psi\rangle$ and the bra $\langle\psi|$ of a given state to yield $\langle\psi|\psi\rangle$. The steps are ordered from top to bottom. In the first step, $C_{[0]}$ is initialized as the 1×1 identity matrix and introduced on the left end of the tensor network, being contracted with the leftmost local tensors $A_{[1]}$ and $A_{[1]}^\dagger$ through the singleton dummy indices β_0 and β'_0 . The contraction of the three tensors $C_{[0]}$, $A_{[1]}$, and $A_{[1]}^\dagger$ —following the strategy described in Fig. 9b—produces the rank-2 tensor $C_{[1]}$. This three-tensor contraction is repeated $N - 1$ times until arriving at the final 1×1 $C_{[N]}$, which is just the desired $\langle\psi|\psi\rangle$. Although this figure considers the computation of the norm of a state $|\psi\rangle$, this scheme can be identically applied to compute the overlap between two distinct MPSs. The closing-the-zipper method can be similarly performed from right to left instead. Assuming the free indices of the MPSs have dimension d and the bond dimension cutoff is D , the closing-the-zipper method cost scales as $\mathcal{O}(ND^3d)$

tensor contraction scheme discussed in Sect. 4.1 (see Fig. 9b), resulting in a $\mathcal{O}(ND^3d) \sim \mathcal{O}(ND^3)$ scaling overall. Unlike the naïf approach, the closing-the-zipper strategy allows for a scalable computation of overlaps between MPSs, which is crucial for the practicality of MPS-based DMRG.


4.3.2 Canonical forms

It is possible to cast the MPS in a suitable form that effectively renders most or even all steps of the closing-the-zipper scheme trivial, thus allowing to simplify the tensor-network diagrams considerably without requiring any detailed calculations. Suppose the MPS is in *left-canonical* form, in which case all local tensors $\{A_{[i]}\}_{i=1}^N$ are left-normalized, that is


or $\sum_{\beta_{i-1}, \sigma_i} A_{\beta'_i, \sigma_i, \beta_{i-1}}^* A_{\beta_{i-1}, \sigma_i, \beta_i} = \delta_{\beta'_i, \beta_i}$ algebraically. In such case, all $\{C_{[n]}\}_{n=0}^N$ in the closing-the-zipper scheme of Fig. 13 are just resolutions of the identity, so all steps are trivial and the MPS is normalized, $\langle\psi|\psi\rangle = 1$. The same conclusions hold if the closing-the-zipper scheme is performed from right to left and the local tensors are all right-normalized

or $\sum_{\beta_i, \sigma_i} A_{\beta_{i-1}, \sigma_i, \beta_i} A_{\beta_i, \sigma_i, \beta'_{i-1}}^* = \delta_{\beta_{i-1}, \beta'_{i-1}}$ algebraically. This is the *right-canonical* form.

For the purposes of computing expectation values of local operators, it is convenient to introduce another canonical form, the so-called *mixed-canonical* MPS, whereby all local tensors to the left of the site on which the local operator acts nontrivially are left-normalized, and all local tensors to the right are right-normalized. To show the usefulness of the mixed-canonical form, let us consider the expectation value $\langle\psi|\hat{O}_{[i]}|\psi\rangle$ of a one-site operator (acting on a given site i) $\hat{O}_{[i]} = \sum_{\sigma_i, \sigma'_i} O_{\sigma_i, \sigma'_i} |\sigma_i\rangle\langle\sigma'_i|$, represented diagrammatically as



Upon making use of the left- and right-normalization of the tensors to the left and to the right of site i , closing the zipper on either side reduces this one-site-operator expectation value to



Any MPS can be converted into left-canonical form by performing SVD on one site at a time, covering the full chain from left to right. As discussed in the final paragraph of Sect. 4.2, each local rank-3 tensor $A_{[i]}$ can be reshaped by fusing the leftmost index with the physical index; the SVD of the resulting matrix yields a unitary matrix U , which becomes a left-normalized rank-3 tensor upon splitting the two indices that were originally fused. Hence, at each site, we replace the original local tensor $A_{[i]}$ with the reshaped U , absorbing the remaining SV^\dagger in the following local tensor $A_{[i+1]}$. At the very last site, because the rightmost index is a singleton dummy index, V^\dagger is just a complex number of modulus 1, so we can neglect it as any wave function is defined up to a global phase factor. Moreover, S is a positive real number that corresponds to the norm of the original MPS. Typically, S is also discarded, in which case the left-canonical MPS becomes normalized. To obtain a right-canonical MPS, one proceeds analogously to the left-canonical case, with the main differences being that the chain is covered from right to left, the local tensor $A_{[i]}$ is replaced by the right-normalized V^\dagger resulting from the SVD at that site, and the remaining US is absorbed by $A_{[i-1]}$. For a mixed-canonical MPS, each of the two processes is carried out on the corresponding side of the selected site.

4.3.3 General wave function representation

Being 1D tensor networks, MPSs are most naturally suited for the representation of wave functions of 1D quantum systems. However, it should be stressed that any wave function, regardless of its dimensionality or entanglement structure, can be represented as an MPS, though possibly with exceedingly large bond dimensions. Suppose we are given the wave function of a quantum system defined on a N -site lattice

$$|\psi\rangle = \sum_{\sigma_1, \sigma_2, \dots, \sigma_N} \psi_{\sigma_1, \sigma_2, \dots, \sigma_N} |\sigma_1\rangle \otimes |\sigma_2\rangle \otimes \dots \otimes |\sigma_N\rangle, \tag{15}$$

where $|\sigma_i\rangle$ denotes the local basis of site i . Assuming the dimension of the local Hilbert space at each site is d , the amplitudes of the wave function, $\psi_{\sigma_1, \sigma_2, \dots, \sigma_N}$, typically cast in the form of a d^N -dimensional vector, can be reshaped into a rank- N tensor such as the one shown in Fig. 8f, with each index having dimension d . To convert this rank- N tensor into the corresponding N -site MPS (Fig. 8e), one can perform SVD at each site at a time following some path that covers every lattice site once.⁹ At the first site, the original rank- N tensor is reshaped into a $d \times d^{N-1}$ matrix; its SVD produces a unitary $d \times d$ matrix U , which is the first local tensor $A_{[1]}$ of the MPS. The remainder of the SVD, the $d \times d^{N-1}$ matrix SV^\dagger , is reshaped into a $d^2 \times d^{N-2}$ matrix, the SVD of which yields a unitary $d^2 \times d^2$ matrix U , which is reshaped into a left-normalized rank-3 tensor with shape $d \times d \times d^2$, corresponding to the second local tensor $A_{[2]}$ of the MPS. This sequence of sitewise SVDs is carried out until reaching the last site, where one obtains a rank-2 tensor $A_{[N]}$ of dimensions $d \times d$. The final outcome is, therefore, a left-canonical MPS. Importantly, because no truncations were performed, until the center of the MPS is reached, the bond dimension keeps on increasing by a factor of d at each site, yielding a maximum bond dimension of $d^{\lfloor N/2 \rfloor}$, which is exponentially large in the system size. This is consistent with the fact that no information was lost, so the number of entries of the MPS is $\mathcal{O}(d^N)$, as for the original rank- N tensor.

The exact conversion of a wave function into an MPS ultimately defeats the purpose of using MPSs (or tensor networks, more generally), which is to provide a more compact representation without compromising the quantitative description of the system under study. A more scalable approach would involve truncating the bond dimension of the MPS to a cutoff D set before-


⁹ In one dimension, the natural choice of path is just to go through the chain from one end to the other. In two and higher dimensions, one may consider following a zigzag path that covers one line of each dimension at a time (see Fig. 4 for an example in two dimensions). In any case, this conversion of a general rank- N tensor into an N -site MPS via a sequence of SVDs works irrespective of the sequence of sites chosen.

hand, although this only produces an approximation of the original state, in general. The remarkable success of MPS-based methods in the study of 1D quantum phenomena is rooted upon the favorable scaling of the required bond dimension cutoff D of MPSs with the system size N , in accordance with the entanglement area laws discussed in Sect. 3.1.3. The relation between the entanglement entropy of a state in a given bipartition and the corresponding bond dimension D of its MPS representation will be clarified in Sect. 5.2.

4.4 Matrix product operators

A *matrix product operator* (MPO) is a 1D tensor network of the form shown diagrammatically in Fig. 14a. The structure of an MPO is similar to that of an MPS, except for the number of physical indices. While an MPS has a single physical index per site, an MPO has two, the top one to act on kets and the bottom one to act on bras, following the convention adopted in Fig. 12. MPOs constitute the most convenient representation of operators for MPS-based methods, as they allow for a sitewise update of the MPS ansatz. In particular, the MPS-based formulation of DMRG discussed in Sect. 5 involves expressing the Hamiltonian under study as an MPO.

Applying an MPO onto an MPS yields another MPS of greater bond dimensions (Fig. 14b). To obtain this MPS, at every site $i = 1, 2, \dots, N$, one contracts the local tensor $A_{[i]}$ from the original MPS with the corresponding local tensor $O_{[i]}$ from the MPO, fusing the pairs of bonds on either side to retrieve a rank-3 tensor $B_{[i]}$



Due to this fusion of indices, the bond dimensions of the final MPS are the product of the bond dimensions of the original MPS and the MPO. The cost of contracting an MPO of bond dimension w (which is typically a small constant, as we shall see below for the case of a short-ranged Hamiltonian) with an MPS of bond dimension D , both with N sites and physical index dimension d , is $\mathcal{O}(ND^2w^2d^2) \sim \mathcal{O}(ND^2)$.

The expectation value of an operator \hat{O} cast in the form of an N -site MPO with respect to a state $|\psi\rangle$ expressed as a N -site MPS is represented in Fig. 14c. One way to obtain $\langle\psi|\hat{O}|\psi\rangle$ is to calculate the MPS corresponding to $\hat{O}|\psi\rangle$ —following the prescription provided in the previous paragraph—and then compute the overlap between the two resulting MPSs, one for $\hat{O}|\psi\rangle$ and another for $\langle\psi|$, using the closing-the-zipper





Fig. 14 **a** Diagram of N -site operator \hat{O} as an MPO. **b** Applying MPO of \hat{O} onto MPS of $|\psi\rangle$ yields another MPS. At m^{th} bond, contracted index β_m of MPS of $\hat{O}|\psi\rangle$ results from fusion of corresponding indices γ_m and α_m of MPO and original MPS, so the dimension of β_m is the product of the dimensions of α_m and γ_m , hence the bold line representation in the diagram. **c** Diagram of expectation value $\langle\psi|\hat{O}|\psi\rangle$. The contraction of the tensor network can be done in two ways. Either the MPO is applied to one of the MPSs, and then, the closing-the-zipper strategy (Fig. 13) is adopted to compute the overlap between the two remaining MPSs or the closing-the-zipper method is applied directly to this three-layer tensor network, as described in the text. Assuming the physical indices have dimension d and the bond dimension cutoffs are D for the MPS and w for the MPO, the cost of the two strategies scales as $\mathcal{O}(ND^3w^2d)$ and $\mathcal{O}(ND^3wd)$, respectively

method introduced in Sect. 4.3.1. The cost of this approach scales as $\mathcal{O}(ND^3w^2d) \sim \mathcal{O}(ND^3)$. Alternatively, the closing-the-zipper strategy can be adapted to contract this three-layer tensor network. Specifically, at the n^{th} iteration, we have



and the indices are contracted as follows:

1. Sum over α_{n-1} with fixed $\sigma_n, \alpha_n, \gamma_{n-1}$ and α'_{n-1} at cost $\mathcal{O}(D^3wd)$.
2. Sum over γ_{n-1} and σ_n with fixed $\alpha_n, \gamma_n, \sigma'_n$ and α'_{n-1} at cost $\mathcal{O}(D^2w^2d^2)$.
3. Sum over α'_{n-1} and σ'_n with fixed α_n, γ_n and α'_n at cost $\mathcal{O}(D^3wd)$.

Upon completing the N iterations to go through all sites, the overall scaling is $\mathcal{O}(ND^3wd) \sim \mathcal{O}(ND^3)$. For technical reasons that will be apparent in Sect. 5.1, this contraction scheme is preferred in the implementation of the finite-system DMRG algorithm.

Any N -site operator can be expressed as an MPO by performing SVD at each site at a time, in a similar spirit to the representation of an arbitrary wave function in terms of an MPS, discussed in Sect. 4.3.3. The problem with this approach is that the bond dimension of the resulting MPO grows by d^2 at every iteration until reaching the middle of the MPO, thus leading to $\mathcal{O}(d^N)$ bond dimensions. The MPO representation of an arbitrary tensor product of single-site operators is straightforward: each local operator is reshaped into a rank-4 tensor with two singleton dummy indices (corresponding to the trivial bonds with dimension $w = 1$), which are contracted with those from the adjacent sites to form the MPO. MPOs like those described above can also be summed¹⁰ to obtain the MPO representation of more generic operators. It must be noted, however, that the previous strategy, although versatile, does not always lead to the lowest possible bond dimensions of

¹⁰ See, e.g., Ref. [53] for a general prescription, which amounts to writing each rank-4 local tensor of the MPOs that we want to sum as a matrix of the physical operators, and then perform direct sums of these matrices at every site, except for the leftmost/rightmost site where the physical operators are organized in a line/column vector.

the final MPO. In particular, it is possible to represent local Hamiltonians in terms of MPOs with $\mathcal{O}(1)$ bond dimension—i.e., constant with respect to the system size N —, as explained below.

The exact MPO of a local Hamiltonian can be obtained through an analytical method originally proposed by McCulloch [21]. For concreteness, let us consider the Heisenberg model for an open-ended spin- s chain with a Zeeman term

$$\hat{\mathcal{H}} = J \sum_{i=1}^{N-1} \hat{\vec{S}}_i \cdot \hat{\vec{S}}_{i+1} - h \sum_{i=1}^N \hat{S}_i^z \tag{16}$$

$$= J \sum_{i=1}^{N-1} \left(\hat{S}_i^z \hat{S}_{i+1}^z + \frac{\hat{S}_i^+ \hat{S}_{i+1}^- + \hat{S}_i^- \hat{S}_{i+1}^+}{2} \right) - h \sum_{i=1}^N \hat{S}_i^z, \tag{17}$$

where J and h are model parameters, $\hat{\vec{S}}_i = (\hat{S}_i^x, \hat{S}_i^y, \hat{S}_i^z)$ is the vector of spin- s operators at site $i \in \{1, 2, \dots, N\}$, and $\hat{S}_i^\pm = \hat{S}_i^x \pm i\hat{S}_i^y$ are the corresponding spin ladder operators. Our goal is to obtain the local tensors $\{H_{[i]}\}_{i=1}^N$ of the MPO that encodes this Hamiltonian. Four different types of terms arise in Eq. (17)

$$\begin{aligned} & \dots \overset{5}{\otimes} \overset{5}{\hat{\mathbb{1}}} \overset{5}{\otimes} J \hat{S}^z \overset{2}{\otimes} \hat{S}^z \overset{1}{\otimes} \overset{1}{\hat{\mathbb{1}}} \overset{1}{\otimes} \dots \\ & \dots \overset{5}{\otimes} \overset{5}{\hat{\mathbb{1}}} \overset{5}{\otimes} \frac{J}{2} \hat{S}^+ \overset{3}{\otimes} \hat{S}^- \overset{1}{\otimes} \overset{1}{\hat{\mathbb{1}}} \overset{1}{\otimes} \dots \\ & \dots \overset{5}{\otimes} \overset{5}{\hat{\mathbb{1}}} \overset{5}{\otimes} \frac{J}{2} \hat{S}^- \overset{4}{\otimes} \hat{S}^+ \overset{1}{\otimes} \overset{1}{\hat{\mathbb{1}}} \overset{1}{\otimes} \dots \\ & \dots \overset{5}{\otimes} \overset{5}{\hat{\mathbb{1}}} \overset{5}{\otimes} -h \hat{S}^z \overset{1}{\otimes} \overset{1}{\hat{\mathbb{1}}} \overset{1}{\otimes} \overset{1}{\hat{\mathbb{1}}} \overset{1}{\otimes} \dots \end{aligned}$$

The numbers above the tensor product signs identify one of the following five ‘states’:

- ‘State’ 1: Only identity operators $\hat{\mathbb{1}}$ to the right.
- ‘State’ 2: One \hat{S}^z operator just to the right, followed by $\hat{\mathbb{1}}$ operators.
- ‘State’ 3: One \hat{S}^- operator just to the right, followed by $\hat{\mathbb{1}}$ operators.
- ‘State’ 4: One \hat{S}^+ operator just to the right, followed by $\hat{\mathbb{1}}$ operators.
- ‘State’ 5: One complete term somewhere to the right.

For a given bulk site i , the local rank-4 tensor $H_{[i]}$, cast in the form of a $w \times w$ matrix where each entry is itself a $d \times d$ matrix—with $w = 5$ the bond dimension of the MPO (determined by the number of ‘states’) and $d = 2s + 1$ the physical index dimension—is constructed in such a way that its (k, l) entry corresponds to the operator that makes the transition from ‘state’ l to ‘state’ k toward the left

$$H_{[i]} = \begin{pmatrix} \hat{1}_i & 0 & 0 & 0 & 0 \\ \hat{S}_i^z & 0 & 0 & 0 & 0 \\ \hat{S}_i^- & 0 & 0 & 0 & 0 \\ \hat{S}_i^+ & 0 & 0 & 0 & 0 \\ -h\hat{S}_i^z & J\hat{S}_i^z & \frac{J}{2}\hat{S}_i^+ & \frac{J}{2}\hat{S}_i^- & \hat{1}_i \end{pmatrix}. \quad (18)$$

For the terminal sites, due to the open boundary conditions, we have two rank-3 tensors, one corresponding to the last row of Eq. (18) for the leftmost site $i = 1$ and another corresponding to the first column of Eq. (18) for the rightmost site $i = N$.

To confirm that the constructed MPO does indeed give rise to the Hamiltonian stated in Eq. (17), one can perform by hand the matrix multiplication of the local tensors in the form shown in Eq. (18), but with the usual scalar multiplications being replaced by tensor products as each entry is itself a rank-2 tensor [54]. Alternatively, the MPO can be contracted and compared directly to the full $d^N \times d^N$ matrix representation of the model Hamiltonian. This sanity check is performed for small system sizes N in the code that complements this manuscript (see Supplementary Information). In this code, we also construct the MPO Hamiltonian for two other quantum spin models, the Majumdar–Ghosh [55, 56] and the Affleck–Kennedy–Lieb–Tasaki [57] models. These two additional examples suffice to demonstrate how to apply McCulloch’s method in general, namely by adding next-nearest-neighbor interactions and further nearest-neighbor interactions, respectively. Assuming the most conventional case of model Hamiltonians with terms acting nontrivially on one or two sites only, the bond dimension of the MPO obtained with this method starts at two and increases by one for every new type of two-site term and/or unit of interaction range [16]. There are, however, notable exceptions to this rule, such as long-range Hamiltonians that allow for a more compact but still exact MPO representation [58, 59]. More complex Hamiltonians such as those arising in quantum chemistry [60] or in two-dimensional lattice models on a cylinder in hybrid real and momentum space [61] may require more sophisticated numerical approaches to reduce the bond dimension of the corresponding MPO [53].

5 Finite-system DMRG in the language of tensor networks

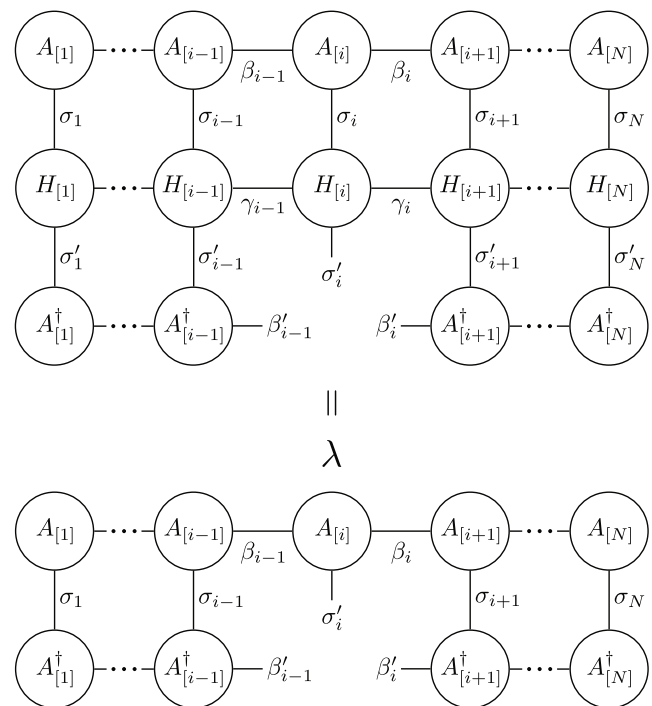
5.1 Derivation: one-site update

The starting point for the derivation of the MPS-based finite-system DMRG algorithm is to consider the set of all N -site MPS representations of a ket $|\psi\rangle$ with (maximum) bond dimension D as a variational space. The local tensor of the MPS at site i is denoted by $A_{[i]}$; for the sake of simplicity, we consider that the physical index dimension is d at all sites. We assume we are given the N -site MPO representation of the

Hamiltonian $\hat{\mathcal{H}}$, with bond dimension $w \sim \mathcal{O}(1)$ and physical index dimension d ; its local rank-4 tensor at site i is denoted by $H_{[i]}$. The goal is to minimize the energy $\langle\psi|\hat{\mathcal{H}}|\psi\rangle$, subject to the normalization constraint $\langle\psi|\psi\rangle = 1$. This can be achieved by minimizing the cost function $\langle\psi|\hat{\mathcal{H}}|\psi\rangle - \lambda\langle\psi|\psi\rangle$, where λ denotes the Lagrange multiplier. The one-site-update version of the algorithm consists of finding the stationary points of the cost function with respect to each local tensor $A_{[i]}^\dagger$ at a time, that is

$$\frac{\partial}{\partial A_{[i]}^\dagger} \left(\langle\psi|\hat{\mathcal{H}}|\psi\rangle - \lambda\langle\psi|\psi\rangle \right) = 0. \quad (19)$$

Making use of the diagrammatic representation, and taking into account that all contractions on a tensor network are linear operations, the derivative with respect to $A_{[i]}^\dagger$ amounts to *punching a hole* [52] at the position of the tensor $A_{[i]}^\dagger$, leading to



which can be understood as a generalized eigenvalue problem for $A_{[i]}$. By casting the MPS in mixed-canonical form with respect to site i , the bottom part of the previous equation simplifies trivially, yielding an eigenvalue problem for $A_{[i]}$ that we write as

$$\sum_a M_{[i]}^{a',a} A_{[i]}^a = \lambda A_{[i]}^{a'}, \quad (20)$$

with $a \equiv (\beta_{i-1}, \sigma_i, \beta_i)$ and $M_{[i]}^{a',a}$ defined by the diagram shown in Fig. 15.

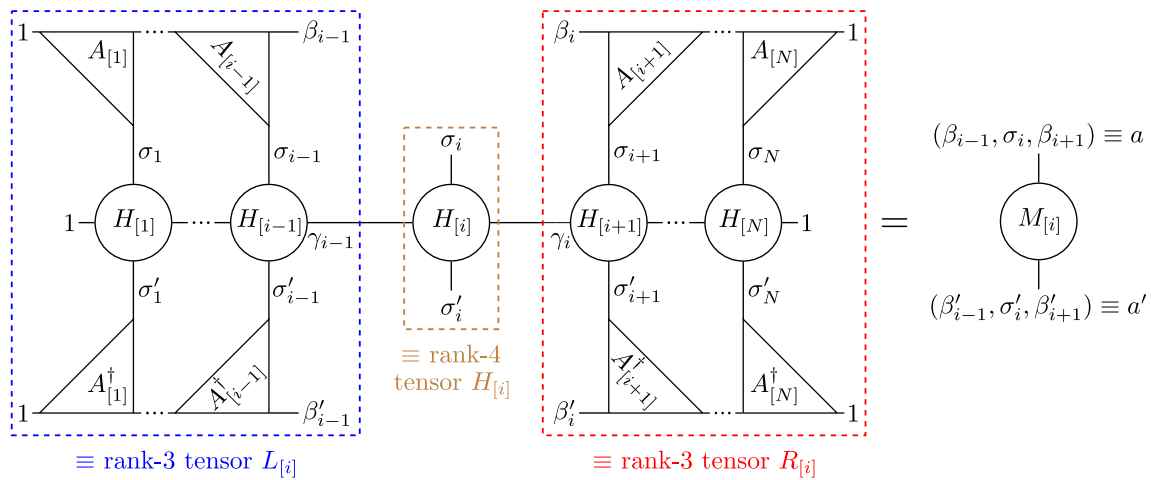


Fig. 15 Tensor-network diagram of the “effective” matrix $M_{[i]}$ of the eigenvalue problem (Eq. (20)) associated with one iteration—corresponding to the local optimization of the MPS at site i —of the one-site-update finite-system DMRG algorithm. To optimize the computational performance of the DMRG algorithm, $M_{[i]}$ is stored in terms of three tensors, $L_{[i]}$, $H_{[i]}$, and $R_{[i]}$ (see Appendix B for details). The pseudocode of the algorithm is shown in Fig. 19 of Appendix D

Having derived an eigenvalue problem [Eq. (20) from the local optimization of the MPS at site i (Eq. (19)], the optimal update of the corresponding local tensor $A_{[i]}$ is simply the eigenstate with lowest eigenvalue, both of which can be obtained through the Lanczos algorithm [62]. Although it is common practice to provide a randomly generated state as the input state to the Lanczos algorithm, in this case, it is preferable to use the current version of the local tensor $A_{[i]}$ as the initial state, as this is a more educated guess of the ground state, thus reducing the number of iterations of the Lanczos algorithm. In addition to the obtained eigenstate being the variationally optimized $A_{[i]}$, the corresponding eigenvalue is also the current estimate of the ground-state energy of the full system. This step of the DMRG algorithm is repeated, sweeping i back and forth between 1 and N . As for the initialization, the typical approach is to start with a random MPS.

Two additional technical remarks regarding the implementation of the DMRG algorithm derived above are in order. First, at every step of the algorithm, after having obtained the updated local tensor $A_{[i]}$ as the ground state of the eigenvalue problem, its SVD is performed to ensure that the MPS is in the appropriate mixed-canonical form in the next step of the sweep, thus avoiding a generalized eigenvalue equation. Second, the “effective” matrix of the eigenvalue problem, $M_{[i]}$, is stored in terms of three separate tensors, $L_{[i]}$, $H_{[i]}$, and $R_{[i]}$ (Fig. 15). As the notation suggests, the rank-4 tensor $H_{[i]}$ is just the local tensor at site i of the MPO that encodes the Hamiltonian $\hat{\mathcal{H}}$. As for the rank-3 tensors $L_{[i]}$ and $R_{[i]}$, they result from the contraction of all tensors to the left and to the right of site i , respectively. The efficient computation of $L_{[i]}$ and $R_{[i]}$ over the multiple sweeps of the DMRG algorithm is detailed in Appendix B.

Making use of the internal structure of the matrix $M_{[i]}$, the time complexity of solving the eigenvalue

problem stated in Eq. (20)—required to update one local tensor of the MPS—is $\mathcal{O}(D^3)$. This scaling results largely from the the matrix-vector multiplications involved in the construction of the Krylov space within the Lanczos algorithm [63]. Note that the naïf explicit contraction of $M_{[i]}$ into a $(D^2d) \times (D^2d)$ matrix would have resulted in a $\mathcal{O}(D^4)$ scaling of the matrix-vector multiplications, as opposed to the $\mathcal{O}(D^3)$ obtained using the $L_{[i]}$, $H_{[i]}$, and $R_{[i]}$ tensors. In the end, all key steps of one iteration of the one-site-update finite-system DMRG algorithm—closing-the-zipper contraction (as described in Appendix B), eigenvalue problem, and SVD—have the same $\mathcal{O}(D^3)$ computational cost, so the overall cost of a full sweep scales as $\mathcal{O}(ND^3)$. It must be noted that, since the standard Python functions to implement the Lanczos algorithm (e.g., `scipy.sparse.linalg.eigsh`) require a matrix as input, the naïf explicit contraction of $M_{[i]}$ was adopted in the code that complements this manuscript, trading efficiency for simplicity.

Finally, although this discussion has been restricted to the computation of the ground state, it is straightforward to extend it to the calculation of low-lying excited states. For concreteness, let us suppose we have already determined the ground state $|\text{GS}\rangle$ in a previous run of the DMRG algorithm and wish to obtain the first excited state. Exploiting the orthogonality of the eigenbasis of the Hamiltonian, we merely have to impose the additional constraint $\langle \psi | \text{GS} \rangle = 0$ through another Lagrange multiplier in the cost function. This additional term effectively imposes an energy penalty on the variational states $|\psi\rangle$ that have nonzero overlap with $|\text{GS}\rangle$. In other words, the eigenvalue problem is restricted to a subspace orthogonal to $|\text{GS}\rangle$. In practice, this condition can be imposed by setting $|\text{GS}\rangle$ as the first Krylov state in the Lanczos algorithm but performing the diagonalization of the tridiagonal matrix

defined in the Krylov subspace spanned by all but the first Krylov state [63].

5.2 Connection to original formalism

The one-site-update MPS-based DMRG algorithm derived in the previous section is entirely analogous to the original formulation of the finite-system DMRG scheme (recall Sect. 3.2) provided that there is only one site—denoted by \circ —instead of two, between the blocks S and E (adopting the notation employed in Fig. 7). Considering a left-to-right sweep¹¹ of the MPS-based formulation, the SVD of the optimized local tensor $A_{[i]}$ at the site i between S and E—which leaves a left-normalized tensor at site i in the MPS representation of the target eigenstate $|\psi\rangle$ —and the subsequent contractions to update the $L_{[i+1]}$ tensor (as defined in Fig. 15) correspond to the projection of the Hilbert space of the growing block S_{\circ} onto the subspace spanned by the highest-weight eigenstates of the reduced density matrix $\sigma_{S_{\circ}} = \text{Tr}_E(|\psi\rangle\langle\psi|)$ considered in the original formulation. The number of kept eigenstates of the reduced density matrix $\sigma_{S_{\circ}}$ in the original formulation is precisely the number of kept singular values in the SVD of the optimized local tensor in the MPS-based version, which translates into the bond dimension D of the MPS ansatz. To support the previous claims, we note that the eigenvalues of $\sigma_{S_{\circ}}$ in the original formulation are the square of the singular values $\{s_n\}_{n=1}^D$ of the SVD of the updated $A_{[i]}$ in the MPS-based version (see Appendix C for the derivation)


$$\sigma_{S_{\circ}} = \text{Tr}_E(|\psi\rangle\langle\psi|) = \sum_{n=1}^D s_n^2 |u_n\rangle_{S_{\circ}} \langle u_n|, \quad (21)$$

where $\{|u_n\rangle_{S_{\circ}}\}_{n=1}^D$ are the D —out of the total Dd —eigenstates of $\sigma_{S_{\circ}}$ with (possibly) nonzero eigenvalues. Therefore, we see that in both formulations of DMRG, D quantifies the degree of entanglement that can be captured across the bipartition between S_{\circ} and E. Moreover, it becomes apparent that the truncation prescribed in the original formulation of the one-site-update finite-system DMRG scheme is actually trivial (see Appendix C for a more detailed explanation), thus sorting out the apparent contradiction related with the fact that no truncation is prescribed in the MPS-based version of the one-site-update DMRG algorithm.

Although the original and the MPS-based formulations of DMRG are equivalent, there is one key difference between them regarding the encoding of the Hamiltonian. While, in the original method, the Hamiltonian obtained from the prior implementation of infinite-system DMRG is inherently approximate as its matrix representation results from an explicit truncation of the Hilbert space through a projection onto a smaller subspace defined by the highest-weight eigen-

states of the reduced density matrices on either side of the bipartition considered, in the MPS-based version, the MPO representation of the Hamiltonian with which one begins to perform the sweeps is exact and the approximate description of the system is entirely restricted to the ansatz of the variational problem, an MPS with given bond dimension D . This difference renders the MPS-based formulation of DMRG more effective at calculating physical quantities related to powers of the Hamiltonian, such as the energy variance or, more generally, cumulant expansions [53].

Similarly to the original formulation of the finite-system DMRG scheme described in Sect. 3.2, there is a two-site-update version of MPS-based DMRG that results from simultaneously minimizing the cost function $\langle\psi|\hat{\mathcal{H}}|\psi\rangle - \lambda\langle\psi|\psi\rangle$ (recall Eq. (19) and the corresponding derivation) with respect to two adjacent local tensors $A_{[i]}^{\dagger}$ and $A_{[i+1]}^{\dagger}$, giving rise to an eigenvalue problem for a two-site tensor of the form



where the explicit contraction (and index fusion) that casts the two-site tensor in the vectorial form C is not carried out in practice (as discussed in Sect. 5.1), but is, nonetheless, a useful picture to have in mind. For the sake of clarity, the labels associated with the legs represent the dimensions of the corresponding indices. Once the eigenvalue problem is solved, the updated tensor C is reshaped into a $(Dd) \times (Dd)$ matrix, so that its SVD can be performed to obtain the optimized local tensors at sites i and $i + 1$. Crucially, the MPS bond dimension between the local tensors at sites i and $i + 1$ increases from D to Dd after this optimization process, so an explicit truncation that keeps only the D highest singular values is required. Hence, the two-site-update algorithm effectively surveys a larger search space than the one-site-update scheme. In particular, this allows to escape local minima in the optimization landscape, namely by having the possibility to explore different symmetry sectors. This is the main reason why the two-site-update DMRG scheme, both in its original and MPS-based formulations, is the standard option in the literature. It should be stressed, however, that the one-site-update DMRG algorithm can be just as reliable as the two-site-update scheme at a lower computational cost (see Sect. 5.4).

In the original formulation of DMRG, the outcome of the infinite-system version is the natural starting point for finite-system DMRG. In the MPS-based version, however, it is common practice to start from a random MPS of given bond dimension D , although this is not usually as good an educated guess as the outcome of the infinite-system DMRG [16]. Alternatively, one may perform finite-system DMRG simulations with MPSs of

¹¹ An analogous reasoning is straightforward for the case of a right-to-left sweep.




Fig. 16 Finite-system DMRG applied to the (isotropic) XY 1D quantum model. Ground-state energy of an open-ended chain composed of 20 $s = 1/2$ spins, as a function of the number of sweeps of the one-site-update DMRG routine, for different values of the bond dimension cutoff D . The dashed black line marks the analytical result [66]. The deviation from the exact solution is also shown in the inset

progressively larger bond dimension D , using the outcome of the previous simulation as the initial state for the current one, padding the remainder of the local tensors with zeros due to the larger bond dimension. A particularly elegant aspect of MPS-based DMRG, notably in its one-site-update scheme, is that the manifold of states explored in the variational problem corresponds to all MPSs of fixed bond dimension D , with no truncations being performed throughout the computation. The tangent-space methods developed in recent years [64, 65] explore this feature in more sophisticated ways.

5.3 Code implementation

In the same spirit of Sect. 3.1.4, we provide a practical implementation of an MPS-based DMRG algorithm. Our code, available both in Supplementary Information and at https://github.com/GCatarina/DMRG_MPS_didactic, consists of a documented Jupyter notebook, written in Python, that goes through all the key steps required to implement the finite-system DMRG method. For simplicity, we consider the one-site update version (see Sect. 5.1), targeting only the ground-state properties. It must also be noted that the coded DMRG routine is model-agnostic, requiring only as input an Hamiltonian MPO with trivial leftmost and rightmost legs. In general, the previous requirement is naturally fulfilled for systems with open boundary conditions in at least one of their physical dimensions.

To benchmark our DMRG code, we apply it to 1D systems, with open boundary conditions, for which exact ground-state solutions are known. Specifically, we consider the (isotropic) XY [67] and the Majumdar–Ghosh [55, 56] models. In Fig. 16, we show how, for

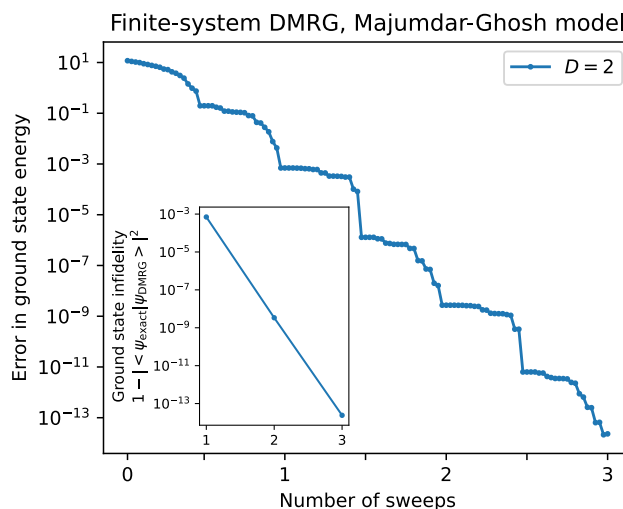


Fig. 17 Finite-system DMRG calculations for the ground state of an open-ended Majumdar–Ghosh chain, whose analytical solution [55, 56] can be written as an MPS with bond dimension 2. Numerical results, obtained with bond dimension cutoff $D = 2$, for a chain composed of 20 $s = 1/2$ spins, show the convergence to the exact ground state as a function of the number of sweeps of the one-site-update DMRG algorithm

an XY chain, the ground-state energy computed by DMRG compares with the analytical result [66]. It is apparent that DMRG converges rapidly with the number of sweeps performed. It is also observed that the accuracy of the numerical calculation is determined by the bond dimension cutoff D . In Fig. 17, a complementary example is shown, where DMRG is used to compute both the error in the energy estimation and the infidelity associated with the ground state of a Majumdar–Ghosh chain. Since, for open-ended Majumdar–Ghosh chains, the exact ground-state wave function (which is unique for chains with even number of spins) can be represented by an MPS with bond dimension 2, it is expected that DMRG yields accurate results with a bond dimension cutoff as small as $D = 2$. It should be noted, however, that, in this case, DMRG takes a few more sweeps to reach convergence.

5.4 State-of-the-art one-site-update DMRG: recent developments

As noted in Sect. 5.2, in the path across the optimization landscape toward the exact ground state, the two-site-update DMRG manages to avoid local minima in which the one-site-update version is often stuck. However, this comes at a cost: the time taken by the two-site-update DMRG to complete a step of a sweep is roughly greater by a factor of d than that of its one-site-update counterpart [16]. Such a difference may not be too relevant for simple examples such as the ones considered in this pedagogical review, but it is certainly significant in state-of-the-art simulations involving large bond dimensions and system sizes (namely

in two-dimensional systems [41]) or even large local Hilbert spaces (e.g., due to high-spin local magnetic moments [68] or multiple fermionic orbitals [69, 70] at each site).

However, it is possible to overcome this limitation of one-site-update DMRG by applying a correction that was first proposed in the language of the original version of DMRG by White in 2005 [71]. The overarching idea consists of adding quantum fluctuations across the bipartition that effectively allow the quantum numbers of the ansatz to change. More concretely, the basis of the environment block E must be enlarged in such a way that the state $\hat{\mathcal{H}}|\psi\rangle$ resulting from applying the Hamiltonian $\hat{\mathcal{H}}$ onto the MPS ansatz $|\psi\rangle$ is contained in the basis. This means that the terms of $\hat{\mathcal{H}}|\psi\rangle$ that connect the system and environment blocks must be added to the density matrix before its reduced form is diagonalized on either side of the bipartition. If the Hamiltonian is split into two parts to reflect the bipartition as¹²


$$\hat{\mathcal{H}} = \sum_{\gamma} h_{\gamma} \hat{A}^{\gamma} \hat{B}^{\gamma}, \tag{22}$$

the correction to the density matrix $\rho = |\psi\rangle\langle\psi|$ for a left-to-right sweep is


$$\Delta\rho = a \sum_{\gamma} \hat{A}^{\gamma} |\psi\rangle\langle\psi| (\hat{A}^{\gamma})^{\dagger}, \tag{23}$$

with $a \sim 10^{-4} - 10^{-3}$ [71] a constant that is a parameter of the simulation. The corrected reduced density matrix on the enlarged system block, $\sigma'_{S_0} = \text{Tr}_E(\rho + \Delta\rho)$, can then be diagonalized to update the basis on which the Hamiltonian is defined. This time, however, contrary to the discussion in Appendix C for the standard one-site-update DMRG, there are, in general, more than D nonzero eigenvalues, so an explicit truncation of the bond dimension is required. This is consistent with the fact that the corrected one-site-update DMRG explores a larger search space, which permits to achieve similar results to those of two-site-update DMRG but at a lower computational cost.

White’s corrected one-site-update DMRG can be straightforwardly formulated in MPS language. Specifically, if the local tensor at the current site i yielded by solving the eigenvalue problem stated in Fig. 15 is $T_{[i]}$, then the analog of $\hat{A}^{\gamma}|\psi\rangle$ in the tensor-network formulation is given by [16]



Notice that the whole portion of the tensor network to the left of site i is just the $L_{[i]}$ tensor used to define the effective matrix in the eigenvalue problem (see Fig. 15). This is already stored in memory, so the determination of the perturbation $P_{[i]}$ is cheap, taking $\mathcal{O}(D^3 dw)$ floating-point operations. However, unlike in standard MPS-based one-site-update DMRG, the reduced density matrix σ'_{S_0} must be constructed explicitly as



where the first term on the right-hand side is the original reduced density matrix σ_{S_0} (see Eq. (21)) and the second term is the correction $\Delta\sigma_{S_0} \equiv \text{Tr}_E \Delta\rho$ [see Eq. (23)]. This extra step is somewhat onerous—taking $\mathcal{O}(D^3 d^2 w)$ floating-point operations—but the corrected single-site-update DMRG is still less costly than the two-site-update version [71]: while the corrected reduced density matrix only has to be computed once for each local tensor update in the former, the two-site-update strategy adopted in the latter adds an extra factor of d to the cost of all K applications of the effective matrix onto the latest Krylov vector in the Lanczos algorithm—as many as the dimension K of the Krylov space required to find the ground state to the desired precision [63].

In 2015, Hubig et al. [72] devised an alternative correction to the one-site-update DMRG that is more suitable for the MPS formulation—in that it forgoes the explicit construction of the reduced density matrix and its diagonalization—and outperforms White’s density-matrix perturbation method [71] in terms of runtime to convergence. For a left-to-right sweep, the MPS ansatz with the updated local tensor $T_{[i]}$ obtained from the eigenvalue problem (see Fig. 15) is

¹² Note that this expression is complete, in the sense that not only does it include the terms of the Hamiltonian that connect the two sides of the bipartition nontrivially, but it also contains the terms that only have support on the system block (for which $B^{\gamma} = \mathbb{1}, \forall \gamma$) and those that only have support on the environment block (for which $A^{\gamma} = \mathbb{1}, \forall \gamma$), assuming a left-to-right sweep.

$$|\psi\rangle = \sum_{\vec{\sigma}} A_{[1]}^{\sigma_1} \dots A_{[i-1]}^{\sigma_{i-1}} T_{[i]}^{\sigma_i} B_{[i+1]}^{\sigma_{i+1}} \dots B_{[N]}^{\sigma_N} |\vec{\sigma}\rangle, \quad (24)$$

where the $\{A_{[j]}^{\sigma_j}\}_{j=1}^{i-1}$ are all left-normalized and the $\{B_{[j]}^{\sigma_j}\}_{j=i+1}^N$ are all right-normalized, as in a mixed-canonical form. $|\psi\rangle$ can be rewritten by performing a subspace expansion at the bond that sets the bipartition

$$\begin{aligned} &\sum_{\vec{\sigma}} A_{[1]}^{\sigma_1} \dots A_{[i-1]}^{\sigma_{i-1}} (T_{[i]}^{\sigma_i} a P_{[i]}^{\sigma_i}) \begin{pmatrix} B_{[i+1]}^{\sigma_{i+1}} \\ 0 \end{pmatrix} \dots B_{[N]}^{\sigma_N} |\vec{\sigma}\rangle \\ &= \sum_{\vec{\sigma}} A_{[1]}^{\sigma_1} \dots A_{[i-1]}^{\sigma_{i-1}} \tilde{T}_{[i]}^{\sigma_i} \tilde{B}_{[i+1]}^{\sigma_{i+1}} \dots B_{[N]}^{\sigma_N} |\vec{\sigma}\rangle. \end{aligned} \quad (25)$$

In words, the local tensor at the current site i is enlarged by adding an expansion term $P_{[i]}$ that effectively allows to probe more environmental states than the original D states. In principle, we are free to choose an arbitrary expansion term, but the one adopted by Hubig et al. [72] is precisely the heuristically motivated $P_{[i]}$ introduced earlier within the context of White’s density-matrix perturbation method, except that it is reshaped into a rank-3 tensor by fusing the two indices of dimensions D and w that define the right-hand-side basis. As a result, the new expanded local tensor $\tilde{T}_{[i]}$ has a $D \times d \times (D + Dw)$ shape; the local tensor to its right, $\tilde{B}_{[i+1]}$, must be expanded accordingly by padding the original tensor with zeros due to the greater bond dimension. At this point, nothing has changed, since Eqs. (24) and (25) are exactly equal. However, the greater bond dimension gives the opportunity to explore a larger search space.

The corrected local tensor can now be obtained without actually referring to the density matrix. Indeed, we can simply do the SVD of the enlarged tensor, $\tilde{T}_{[i]} = USV^\dagger$, truncating the bond dimension back to D . As usual, the reshaped left-normalized U becomes the final update of the local tensor at site i , while SV^\dagger is absorbed in the next tensor at site $i+1$. Of course, the final D states selected by the truncation after the SVD are not necessarily the same as the original D states with nonzero amplitudes, so this effectively allows to escape local minima in the optimization process. All in all, the construction of the expansion term $P_{[i]}$ (detailed in the previous page) and the SVD of the enlarged tensor $\tilde{T}_{[i]} = (T_{[i]} a P_{[i]})$ take $\mathcal{O}(D^3 dw)$ and $\mathcal{O}(D^3 dw^2)$ floating-point operations, respectively, so no steps with scaling $\mathcal{O}(D^3 d^2)$ are present in this method. This justified the name *strictly single-site DMRG* (DMRG3S) coined by Hubig et al. [72].

Thus far, we have not considered the role played by the heuristic parameter a , which basically sets the order of magnitude of the quantum fluctuations introduced in the state. On the one hand, too large a value of a hinders convergence by obscuring the improvements made by the local optimizer. On the other hand, too small a value of a does not allow to avoid the local minima traps. As a result, the value of a must be judi-

ciously tuned throughout the simulation, being model-dependent [73]. Still, even after tuning a , it is possible that the energy of the ansatz is increased upon truncating the bond dimension [72].

This issue was addressed in 2022 by Gleis et al. [73] by devising a fully variational (in the sense that the energy estimate never increases) version of the corrected single-site-update DMRG that converges more rapidly than the subspace-expansion method proposed earlier by Hubig et al. [72]. The method by Gleis et al. [73] is based on a *controlled bond expansion* (CBE), which amounts to identifying parts of the two-site-update orthogonal space that carry significant weight in $\hat{\mathcal{H}}|\psi\rangle$ and to include only those parts when expanding the virtual bonds of the single-site-update Hamiltonian. Remarkably, these parts can be found via a projector that can be constructed at single-site-update costs through a so-called *shrewd selection*. Importantly, the CBE method makes use of no mixing parameters such as a ; there is one simulation parameter that controls the degree of bond expansion, but it was found to be model-independent and to remain constant throughout the calculation [73]. Moreover, the CBE method was shown to generically converge significantly faster with respect to the number of sweeps than the subspace-expansion method [72] while taking about the same CPU time per sweep, thus resulting in an overall significantly faster convergence [73].

On a final note, although these corrections to one-site-update DMRG were not implemented in the code that accompanies this pedagogical review—as they were not relevant for the examples herein considered—we encourage the readers to implement them on their own as an exercise. The implementation of White’s density-matrix perturbation method [71] and of the subspace-expansion method by Hubig et al. [72] should be straightforward; the only potentially tricky technical issue is the update of the mixing factor a , for which Section VI of Ref. [72] can be helpful. The state-of-the-art CBE method by Gleis et al. [73] does involve a slightly more sophisticated enrichment method, but, in the end, it amounts to a series of SVDs that are carefully detailed in the supplemental material of Ref. [73].

6 Conclusion

In summary, we have provided a comprehensive introduction to DMRG, both in the original and in the tensor-network formulations. For pedagogical purposes, our work is accompanied by concrete practical implementations (see Supplementary Information), the main goal of which is to make the formal description of the method more tangible. For that reason, our efforts were directed toward producing digestible, transparent, and instructive code implementations rather than optimizing their performance or versatility. Although there exist publicly available user-friendly libraries that efficiently implement DMRG (e.g., TeNPy [17] and ITen-

sor [18]), we believe that a clear understanding of the method is crucial for an educated use of these resources. Moreover, it is our opinion that the fundamentals of DMRG are interesting in their own right, as they are at the same time powerful and simple.

Despite not having been covered in this colloquium, extensions of DMRG to tackle quantum dynamics [74–77] and finite-temperature behavior [78, 79]—both relevant for the study of out-of-equilibrium quantum many-body phenomena—have been put forth. Another topic that was beyond the scope of this review was the exploitation of symmetries [13, 80, 81] to restrict the DMRG simulations to a given symmetry subspace, both to speed-up the calculations and to find excited states without having to compute and impose the orthogonality with respect to all the lower energy states. In any event, upon completing the reading of this manuscript, we are confident that the reader is ready to explore the relevant literature to become acquainted with these methods.

Acknowledgements G. C. acknowledges financial support from Fundação para a Ciência e a Tecnologia (FCT) for the PhD scholarship grant with reference No. SFRH/BD/138806/2018. B. M. acknowledges financial support from FCT for the PhD scholarship grant with Reference No. SFRH/BD/08444/2020.

Author contributions

GC wrote the code implementations. All authors contributed toward writing the manuscript.

Funding Information Open Access funding provided by Lib4RI – Library for the Research Institutes within the ETH Domain: Eawag, Empa, PSI & WSL.

Data Availability Statement The authors declare that the data supporting this work are available within the manuscript and its supplementary information files.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A: Gauge freedom of SVD

In the factorization of a matrix via singular value decomposition (SVD), the singular values are unique, but not the singular vectors in general [82].

Given a singular value decomposition USV^\dagger , there is a gauge freedom associated with the introduction of a resolution of the identity, $\mathbb{1} = WW^\dagger$, in $US(WW^\dagger)V^\dagger$. Provided that W commutes with S , the matrices W and W^\dagger can be absorbed in the definition of the left- and right-singular vectors to produce an alternative singular value decomposition $\tilde{U}\tilde{S}\tilde{V}^\dagger$, with $\tilde{U} \equiv UW$ and $\tilde{V}^\dagger \equiv W^\dagger V^\dagger$. If all singular values are distinct, then W must also be diagonal to commute with S , in which case left- and right-singular vectors are unique up to a phase factor $e^{i\theta}$. If, instead, there are repeated eigenvalues, then the associated left- and right-singular vectors may be chosen in any fashion, such that they span the relevant subspace. This corresponds to W being a block-diagonal unitary matrix, with nontrivial blocks associated with the singular vectors of equal singular values.

It is also possible to introduce on either side of S two resolutions of the identity constructed with a permutation matrix P , $U(P^\dagger P)S(P^\dagger P)V^\dagger$, and absorb the permutation matrices as in $U'S'V'^\dagger$, with $U' \equiv UP^\dagger$, $V'^\dagger \equiv PV^\dagger$, and $S' \equiv PSP^\dagger$. The matrix S' is still diagonal, but the order of the entries along the diagonal has changed relative to S . In particular, this gauge transformation allows to rearrange the singular values in descending order in the definition of S . This is a common practice, particularly when truncations are considered.

B: Update of “effective” matrix in one-site-update MPS-based DMRG

In this appendix, we explain how to initialize and efficiently update the rank-3 $L_{[i]}$ and $R_{[i]}$ tensors that are part of the structure of the “effective” matrix $M_{[i]}$ (see Fig. 15) of the eigenvalue problem that results in the optimal update of the local tensor $A_{[i]}$ at site $i \in \{1, 2, \dots, N\}$ of the MPS ansatz (with bond dimension D) within the one-site-update finite-system DMRG algorithm. The Hamiltonian considered has an N -site MPO representation with $\mathcal{O}(1)$ bond dimension; its local tensor at site i is denoted by $H_{[i]}$.

Let us assume that the very first site of the MPS to be optimized is the leftmost site, $i = 1$. In that case, the initial MPS should be cast in right-canonical form, which takes $\mathcal{O}(ND^3)$ operations. Before initializing the first left-to-right sweep, a preliminary right-to-left routine (without any local optimization of the MPS) is carried out to compute all $\{R_{[i]}\}_{i=1}^N$ sequentially. The initial $R_{[N]}$ is just the $1 \times 1 \times 1$ (reshaped) identity, and then, $R_{[N-1]}$ is obtained by contracting the right-normalized $A_{[N]}$, its adjoint $A_{[N]}^\dagger$ and $H_{[N]}$ with $R_{[N]}$, following the three-layer closing-the-zipper strategy introduced in Sect. 4.4. At the end of this preliminary right-to-left routine, all $\{R_{[i]}\}_{i=1}^N$ have been computed in $\mathcal{O}(ND^3)$ time. Therefore, we see that the initialization of the DMRG algorithm has a computational cost of $\mathcal{O}(ND^3)$.

At this point, all tensors required to define the eigenvalue problem at the first site are available, since the corresponding $L_{[1]}$ operator is the trivial $1 \times 1 \times 1$ identity—there is

nothing to the left of site $i = 1$. We can therefore start the first left-to-right sweep to optimize the MPS. At the end of a given iteration i of this left-to-right sweep, corresponding to the optimization of the local tensor $A_{[i]}$ at site i , the rank-3 tensor $L_{[i+1]}$ is computed—contracting the previously determined $L_{[i]}$ with the left-normalized $A_{[i]}$, its adjoint $A_{[i]}^\dagger$ and $H_{[i]}$, at $\mathcal{O}(D^3)$ cost—so that it can be used to define the eigenvalue problem of the next iteration. During such left-to-right sweep, the $R_{[i]}$ tensors do not have to be recalculated, because the updated tensors are all absorbed by the $L_{[i]}$ tensors. Importantly, only a single $L_{[i]}$ tensor is computed at every iteration of the sweep, so the time complexity of one iteration is $\mathcal{O}(D^3)$.

Once a left-to-right sweep is completed and we move on to a right-to-left sweep, the roles are reversed: the $L_{[i]}$ tensors are retrieved from memory and the $R_{[i]}$ tensors are recalculated iteratively.

C: Trivial truncation of Hilbert space in original version of one-site-update finite-system DMRG

In the original formulation of the one-site-update finite-system DMRG (see Fig. 7 but consider only one site, denoted by \circ , between blocks S and E), for a d -dimensional local degree of freedom and D kept eigenstates, the Hamiltonian of the full system, $S \circ E$, is a $(D^2d) \times (D^2d)$ matrix, so the target eigenstate $|\psi\rangle$ is a (D^2d) -dimensional vector, which is computed. Assuming a left-to-right sweep, without loss of generality, the full system in the next iteration, which we denote by $S' \circ E'$, has an Hilbert space with increased dimension D^2d^2 , since the $D \times D$ Hamiltonian of the shrunk block E' is fetched from memory, but the Hamiltonian of the grown block $S' \equiv S \circ$ is obtained anew, yielding a $(Dd) \times (Dd)$ matrix. Therefore, the Hilbert space of the block S' has to be truncated before the diagonalization of the Hamiltonian of $S' \circ E'$ takes place.

To compute the reduced density matrices on either side of the bipartition between $S \circ$ and E , it is useful to obtain the Schmidt decomposition [83] of $|\psi\rangle$. This involves reshaping the (D^2d) -dimensional vector $|\psi\rangle$ into a $(Dd) \times D$ matrix M —according to the considered bipartition—and then performing its full SVD (see Sect. 4.2). This yields $M = \mathcal{U}S\mathcal{V}^\dagger$, with \mathcal{U} a $(Dd) \times (Dd)$ unitary matrix with columns $\{|u_n\rangle\}_{n=1}^{Dd}$ (the left-singular vectors), S a $(Dd) \times D$ matrix with non-negative real entries along the diagonal (the singular values $\{s_n\}_{n=1}^D$) and all remaining entries equal to zero, and \mathcal{V}^\dagger a $D \times D$ unitary matrix with lines $\{|v_n\rangle\}_{n=1}^D$ (the right-singular vectors).

However, as noted in Sect. 4.2, by considering the thin SVD instead, it is possible to convert \mathcal{S} , the matrix that encodes the singular values, into a $D \times D$ matrix S by discarding the corresponding $(Dd - D)$ columns of \mathcal{U} , resulting in the left-normalized $(Dd) \times D$ matrix U . These discarded columns are nothing more than left-singular vectors associated with zero-valued rows of \mathcal{S} , so this truncation is exact.

In the end, the Schmidt decomposition reads as

$$|\psi\rangle = \sum_{n=1}^D s_n |u_n\rangle_{S \circ} \otimes |v_n\rangle_E, \tag{26}$$

and the reduced density matrices on either side can be written as

$$\begin{aligned} \sigma_{S \circ} &= \text{Tr}_E(|\psi\rangle\langle\psi|) = \sum_{n=1}^D s_n^2 |u_n\rangle_{S \circ} \langle u_n|, \\ \sigma_E &= \text{Tr}_{S \circ}(|\psi\rangle\langle\psi|) = \sum_{n=1}^D s_n^2 |v_n\rangle_E \langle v_n|. \end{aligned} \tag{27}$$

In summary, we have shown that the eigenvalues of the reduced density matrices are the square of the singular values obtained by performing the SVD of the target eigenstate in the corresponding bipartition, thus establishing a connection between the original and the MPS-based formulations of DMRG. Moreover, we have found that $\sigma_{S \circ}$, which is generally a $(Dd) \times (Dd)$ matrix, only has D eigenvectors with nonzero eigenvalues, which can be used to truncate the Hilbert space of the block S' without any approximation, thus showing why no actual truncation takes place in the original formulation of the one-site update finite-system DMRG algorithm, as in the corresponding MPS-based version.

D: Pseudocodes

In Fig. 18, we present the pseudocode of the infinite-system DMRG algorithm, within the original formulation. The corresponding code implementation is available both in Supplementary Information and at https://github.com/GCatarina/DMRG_didactic.

Figure 19 shows the pseudocode of the MPS-based one-site-update finite-system DMRG algorithm. The corresponding code implementation is available both in Supplementary Information and at https://github.com/GCatarina/DMRG_MPS_didactic.

Algorithm 1: Original infinite-system DMRG

```

// Inputs:
// nr of iterations ( $N_{it}$ ) relates to final size of system
// nr of kept eigenvectors ( $D$ ) relates to precision of
// algorithm
// -----
for  $k=1, \dots, N_{it}$  do
  if  $k=1$  then
    // S
    Build Hamiltonian of S ( $\mathcal{H}_S$ )
    Build relevant operators of S ( $X_S$ ) for interaction between S and l ( $\mathcal{V}_{S-l}$ )
    // E
    Build Hamiltonian of E ( $\mathcal{H}_E$ )
    Build relevant operators of E ( $X_E$ ) for interaction between r and E ( $\mathcal{V}_{r-E}$ )
  else
    // S
     $\mathcal{H}_S = O_L \mathcal{H}_L O_L^\dagger$ 
     $X_S = O_L X_L O_L^\dagger$ 
    // E
     $\mathcal{H}_E = O_R \mathcal{H}_R O_R^\dagger$ 
     $X_E = O_R X_R O_R^\dagger$ 
  end
  //  $L = S \otimes l$ 
  Build Hamiltonian of l ( $\mathcal{H}_l$ )
  Build Hamiltonian of L ( $\mathcal{H}_L = \mathcal{H}_S \otimes \mathbb{1}_l + \mathbb{1}_S \otimes \mathcal{H}_l + \mathcal{V}_{S-l}$ )
  Build relevant operators of L ( $X_L$ ) for interaction between L and R ( $\mathcal{V}_{L-R}$ )
  //  $R = r \otimes E$ 
  Build Hamiltonian of r ( $\mathcal{H}_r$ )
  Build Hamiltonian of R ( $\mathcal{H}_R = \mathcal{H}_r \otimes \mathbb{1}_E + \mathbb{1}_r \otimes \mathcal{H}_E + \mathcal{V}_{r-E}$ )
  Build relevant operators of R ( $X_R$ ) for interaction between L and R
  //  $SB = L \otimes R$ 
  Build Hamiltonian of SB ( $\mathcal{H}_{SB} = \mathcal{H}_L \otimes \mathbb{1}_R + \mathbb{1}_L \otimes \mathcal{H}_R + \mathcal{V}_{L-R}$ )
  Diagonalize  $\mathcal{H}_{SB}$ , obtaining eigenvector  $|\psi\rangle$  with lowest eigenvalue
  // Truncation L
  Build reduced density matrix on L ( $\sigma_L = \text{Tr}_R\{|\psi\rangle\langle\psi|\}$ )
  Diagonalize  $\sigma_L$ 
  Build  $O_L$  where rows contain  $D$  eigenvectors of  $\sigma_L$  with lowest eigenvalue
  // Truncation R
  Build reduced density matrix on R ( $\sigma_R = \text{Tr}_L\{|\psi\rangle\langle\psi|\}$ )
  Diagonalize  $\sigma_R$ 
  Build  $O_R$  where rows contain  $D$  eigenvectors of  $\sigma_R$  with lowest eigenvalue
end

```

Fig. 18 Pseudocode of the original infinite-system DMRG algorithm, as described in the workflow diagram of Fig. 3. For concreteness, this pseudocode assumes that the target state is the ground state. It is also straightforward to generalize this pseudocode to the case of an adaptive algorithm where, instead of fixing D , the truncation error given by Eq. (9) is kept below a given threshold

Algorithm 2: MPS-based one-site-update finite-system DMRG

```

// Inputs:
// Hamiltonian MPO ( $H[i]$ ,  $i = 1, \dots, N$ ) defines model of system
// with  $N$  sites
// maximum bond dimension ( $D$ ) relates to precision of
// algorithm
// nr of sweeps ( $N_{\text{sweep}}$ ) relates to precision of algorithm
// -----
// Initialization
Build random MPS ( $A[i]$ ,  $i = 1, \dots, N$ ) with bond dimension  $D$ 
Transform  $A[i]$  to right-canonical form ▷ SVDs
Set  $L[1] = R[N] = \mathbb{1}$ 
for  $i=N-1, \dots, 1$  do
  | Compute  $R[i]$  ▷ via  $R[i+1], A[i], H[i], A^\dagger[i]$ 
end
// Sweeping
for  $k=1, \dots, N_{\text{sweep}}$  do
  // left-to-right sweep
  for  $i=1, \dots, N$  do
    // local optimization of  $A[i]$ 
    Compute  $M[i]$  ▷ via  $L[i], H[i], R[i]$ 
    Diagonalize  $M[i]$  ▷ Lanczos
     $A[i] \leftarrow$  lowest-energy eigenvector of  $M[i]$ 
    // preparing next iteration
    SVD  $A[i]$ , obtaining  $USV^\dagger$ 
     $A[i] \leftarrow U$ 
    if  $i \neq N$  then
      | Compute  $L[i+1]$  ▷ via  $L[i], A[i], H[i], A^\dagger[i]$ 
    end
  end
  // right-to-left sweep
  for  $i=N, \dots, 1$  do
    // local optimization of  $A[i]$ 
    Compute  $M[i]$  ▷ via  $L[i], H[i], R[i]$ 
    Diagonalize  $M[i]$  ▷ Lanczos
     $A[i] \leftarrow$  lowest-energy eigenvector of  $M[i]$ 
    // preparing next iteration
    SVD  $A[i]$ , obtaining  $USV^\dagger$ 
     $A[i] \leftarrow V^\dagger$ 
    if  $i \neq 1$  then
      | Compute  $R[i-1]$  ▷ via  $R[i], A[i], H[i], A^\dagger[i]$ 
    end
  end
end

```

Fig. 19 Pseudocode of the MPS-based one-site-update finite-system DMRG algorithm, using the notation of Fig. 15. This pseudocode assumes an algorithm that targets the ground-state properties. For the sake of brevity, reshaping operations of tensors are not explicitly indicated

References

1. B. Keimer, J.E. Moore, The physics of quantum materials. *Nat. Phys.* **13**(11), 1045–1055 (2017). <https://doi.org/10.1038/nphys4302>
2. W. Kohn, Nobel lecture: Electronic structure of matter—wave functions and density functionals. *Rev. Mod. Phys.* **71**, 1253–1266 (1999). <https://doi.org/10.1103/RevModPhys.71.1253>
3. C.D. Sherrill, H.F. Schaefer III., *Advances in Quantum Chemistry*, vol. 34 (Academic Press, New York, 1999), pp.143–269. [https://doi.org/10.1016/S0065-3276\(08\)60532-8](https://doi.org/10.1016/S0065-3276(08)60532-8)
4. P. Hohenberg, W. Kohn, Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871 (1964). <https://doi.org/10.1103/PhysRev.136.B864>
5. W. Kohn, L.J. Sham, Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965). <https://doi.org/10.1103/PhysRev.140.A1133>
6. F. Giustino, *Materials Modelling Using Density Functional Theory* (Oxford University Press, Oxford, 2014)
7. W.M.C. Foulkes, L. Mitas, R.J. Needs, G. Rajagopal, Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.* **73**, 33–83 (2001). <https://doi.org/10.1103/RevModPhys.73.33>
8. R.P. Feynman, Simulating physics with computers. *Int. J. Theor. Phys.* **21**(6), 467–488 (1982). <https://doi.org/10.1007/BF02650179>
9. S. Lloyd, Universal quantum simulators. *Science* **273**(5278), 1073–1078 (1996). <https://doi.org/10.1126/science.273.5278.1073>
10. I.M. Georgescu, S. Ashhab, F. Nori, Quantum simulation. *Rev. Mod. Phys.* **86**, 153–185 (2014). <https://doi.org/10.1103/RevModPhys.86.153>
11. S.R. White, Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992). <https://doi.org/10.1103/PhysRevLett.69.2863>
12. S.R. White, Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* **48**, 10345–10356 (1993). <https://doi.org/10.1103/PhysRevB.48.10345>
13. U. Schollwöck, The density-matrix renormalization group. *Rev. Mod. Phys.* **77**, 259–315 (2005). <https://doi.org/10.1103/RevModPhys.77.259>
14. S. Östlund, S. Rommer, Thermodynamic limit of density matrix renormalization. *Phys. Rev. Lett.* **75**, 3537–3540 (1995). <https://doi.org/10.1103/PhysRevLett.75.3537>
15. J. Dukelsky, M.A. Martín-Delgado, T. Nishino, G. Sierra, Equivalence of the variational matrix product method and the density matrix renormalization group applied to spin chains. *EPL (Europhysics Letters)* **43**(4), 457–462 (1998). <https://doi.org/10.1209/epl/i1998-00381-x>
16. U. Schollwöck, The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **326**(1), 96–192 (2011). <https://doi.org/10.1016/j.aop.2010.09.012>
17. J. Hauschild, F. Pollmann, Efficient numerical simulations with tensor networks: Tensor network python (TeNPy). *SciPost Phys. Lect. Notes* p. 5 (2018). <https://doi.org/10.21468/SciPostPhysLectNotes.5>
18. M. Fishman, S.R. White, E.M. Stoudenmire, The ITensor software library for tensor network calculations. *SciPost Phys. Codebases* p. 4 (2022). <https://doi.org/10.21468/SciPostPhysCodeb.4>
19. R.M. Noack, S.R. Manmana, Diagonalization-and numerical renormalization-group-based methods for interacting quantum systems. *AIP Conf. Proc.* **789**(1), 93–163 (2005). <https://doi.org/10.1063/1.2080349>
20. K.A. Hallberg, New trends in density matrix renormalization. *Adv. Phys.* **55**(5–6), 477–526 (2006). <https://doi.org/10.1080/00018730600766432>
21. I.P. McCulloch, From density-matrix renormalization group to matrix product states. *J. Stat. Mech. Theory Exp.* **10**, P10,014 (2007). <https://doi.org/10.1088/1742-5468/2007/10/p10014>
22. G. De Chiara, M. Rizzi, D. Rossini, S. Montangero, Density matrix renormalization group for dummies. *J. Comput. Theor. Nanosci.* **5**(7), 1277–1288 (2008). <https://doi.org/10.1166/jctn.2008.2564>
23. T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, C. Willing, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, ed. by F. Loizides, B. Schmidt (IOS Press, 2016), pp. 87 – 90
24. K.G. Wilson, The renormalization group: critical phenomena and the Kondo problem. *Rev. Mod. Phys.* **47**, 773–840 (1975). <https://doi.org/10.1103/RevModPhys.47.773>
25. R. Bulla, T.A. Costi, T. Pruschke, Numerical renormalization group method for quantum impurity systems. *Rev. Mod. Phys.* **80**, 395–450 (2008). <https://doi.org/10.1103/RevModPhys.80.395>
26. G. Van Rossum, F.L. Drake, *Python 3 Reference Manual* (CreateSpace, Scotts Valley, 2009)
27. J. Kondo, Resistance minimum in dilute magnetic alloys. *Progress Theoret. Phys.* **32**(1), 37–49 (1964). <https://doi.org/10.1143/PTP.32.37>
28. P.W. Anderson, Localized magnetic states in metals. *Phys. Rev.* **124**, 41–53 (1961). <https://doi.org/10.1103/PhysRev.124.41>
29. S.R. White, R.M. Noack, Real-space quantum renormalization groups. *Phys. Rev. Lett.* **68**, 3487–3490 (1992). <https://doi.org/10.1103/PhysRevLett.68.3487>
30. G. Vidal, Classical simulation of infinite-size quantum lattice systems in one spatial dimension. *Phys. Rev. Lett.* **98**, 070,201 (2007). <https://doi.org/10.1103/PhysRevLett.98.070201>
31. I.P. McCulloch, Infinite Size Density Matrix Renormalization Group, Revisited. [arXiv:0804.2509](https://arxiv.org/abs/0804.2509) (2008)
32. V. Zauner-Stauber, L. Vanderstraeten, M.T. Fishman, F. Verstraete, J. Haegeman, Variational optimization algorithms for uniform matrix product states. *Phys. Rev. B* **97**, 045,145 (2018). <https://doi.org/10.1103/PhysRevB.97.045145>
33. I. Schur, Über eine Klasse von Mittelbildungen mit Anwendungen auf die Determinantentheorie. *Sitzungsber. Berl. Math. Ges.* **22**(9–20), 51 (1923)
34. A. Horn, Doubly stochastic matrices and the diagonal of a rotation matrix. *Am. J. Math.* **76**, 620 (1954). <https://doi.org/10.2307/2372705>
35. P. Calabrese, A. Lefevre, Entanglement spectrum in one-dimensional systems. *Phys. Rev. A* **78**, 032,329 (2008). <https://doi.org/10.1103/PhysRevA.78.032329>
36. F. Pollmann, S. Mukerjee, A.M. Turner, J.E. Moore, Theory of finite-entanglement scaling at one-dimensional quantum critical points. *Phys. Rev. Lett.* **102**, 255,701 (2009). <https://doi.org/10.1103/PhysRevLett.102.255701>

37. J. Eisert, M. Cramer, M.B. Plenio, Colloquium: area laws for the entanglement entropy. *Rev. Mod. Phys.* **82**, 277–306 (2010). <https://doi.org/10.1103/RevModPhys.82.277>
38. M.B. Hastings, An area law for one-dimensional quantum systems. *J. Stat. Mech. Theory Exp.* **08**, P08,024 (2007). <https://doi.org/10.1088/1742-5468/2007/08/p08024>
39. G. Vidal, J.I. Latorre, E. Rico, A. Kitaev, Entanglement in quantum critical phenomena. *Phys. Rev. Lett.* **90**, 227,902 (2003). <https://doi.org/10.1103/PhysRevLett.90.227902>
40. J.I. Latorre, E. Rico, G. Vidal, Ground state entanglement in quantum spin chains. *Quantum Inf. Comput.* **4**(1), 48–92 (2004)
41. E.M. Stoudenmire, S.R. White, Studying two-dimensional systems with the density matrix renormalization group. *Ann. Rev. Condens. Matter Phys.* **3**(1), 111–128 (2012). <https://doi.org/10.1146/annurev-conmatphys-020911-125018>
42. S.R. White, D.A. Huse, Numerical renormalization-group study of low-lying eigenstates of the antiferromagnetic $S = 1$ Heisenberg chain. *Phys. Rev. B* **48**, 3844–3852 (1993). <https://doi.org/10.1103/PhysRevB.48.3844>
43. F. Verstraete, D. Porras, J.I. Cirac, Density matrix renormalization group and periodic boundary conditions: a quantum information perspective. *Phys. Rev. Lett.* **93**, 227,205 (2004). <https://doi.org/10.1103/PhysRevLett.93.227205>
44. R. Orús, A practical introduction to tensor networks: matrix product states and projected entangled pair states. *Ann. Phys.* **349**, 117–158 (2014). <https://doi.org/10.1016/j.aop.2014.06.013>
45. J.C. Bridgeman, C.T. Chubb, Hand-waving and interpretive dance: an introductory course on tensor networks. *J. Phys. A Math. Theor.* **50**(22), 223,001 (2017). <https://doi.org/10.1088/1751-8121/aa6dc3>
46. J. Biamonte, V. Bergholm, Tensor Networks in a Nutshell. [arXiv:1708.00006](https://arxiv.org/abs/1708.00006) (2017)
47. L. Chi-Chung, P. Sadayappan, R. Wenger, On optimizing a class of multi-dimensional loops with reduction for parallel execution. *Parallel Process. Lett.* **07**(02), 157–168 (1997). <https://doi.org/10.1142/S0129626497000176>
48. R.N.C. Pfeifer, J. Haegeman, F. Verstraete, Faster identification of optimal contraction sequences for tensor networks. *Phys. Rev. E* **90**, 033,315 (2014). <https://doi.org/10.1103/PhysRevE.90.033315>
49. C. Eckart, G. Young, The approximation of one matrix by another of lower rank. *Psychometrika* **1**(3), 211–218 (1936). <https://doi.org/10.1007/BF02288367>
50. L. Mirsky, Symmetric gauge functions and unitarily invariant norms. *Q. J. Math.* **11**(1), 50–59 (1960). <https://doi.org/10.1093/qmath/11.1.50>
51. G.H. Golub, C.F. Van Loan, *Matrix Computation*, 3rd edn. (The Johns Hopkins University Press, Baltimore, 1996)
52. J. von Delft. Tensor Networks Lecture Course: Physics Department, Ludwig-Maximilians-Universität München (2020). https://www2.physik.uni-muenchen.de/lehre/vorlesungen/sose_20/tensor_networks_20/index.html
53. C. Hubig, I.P. McCulloch, U. Schollwöck, Generic construction of efficient matrix product operators. *Phys. Rev. B* **95**, 035,129 (2017). <https://doi.org/10.1103/PhysRevB.95.035129>
54. G.M. Crosswhite, D. Bacon, Finite automata for caching in matrix product algorithms. *Phys. Rev. A* **78**, 012,356 (2008). <https://doi.org/10.1103/PhysRevA.78.012356>
55. C.K. Majumdar, D.K. Ghosh, On next-nearest-neighbor interaction in linear chain. I. *J. Math. Phys.* **10**(8), 1388–1398 (1969). <https://doi.org/10.1063/1.1664978>
56. C.K. Majumdar, D.K. Ghosh, On next-nearest-neighbor interaction in linear chain. II. *J. Math. Phys.* **10**(8), 1399–1402 (1969). <https://doi.org/10.1063/1.1664979>
57. I. Affleck, T. Kennedy, E.H. Lieb, H. Tasaki, Rigorous results on valence-bond ground states in antiferromagnets. *Phys. Rev. Lett.* **59**, 799–802 (1987). <https://doi.org/10.1103/PhysRevLett.59.799>
58. G.M. Crosswhite, A.C. Doherty, G. Vidal, Applying matrix product operators to model systems with long-range interactions. *Phys. Rev. B* **78**, 035,116 (2008). <https://doi.org/10.1103/PhysRevB.78.035116>
59. F. Fröwis, V. Nebendahl, W. Dür, Tensor operators: constructions and applications for long-range interaction systems. *Phys. Rev. A* **81**, 062,337 (2010). <https://doi.org/10.1103/PhysRevA.81.062337>
60. G.K.L. Chan, S. Sharma, The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.* **62**(1), 465–481 (2011). <https://doi.org/10.1146/annurev-physchem-032210-103338>. (PMID: **21219144**)
61. J. Motruk, M.P. Zaletel, R.S.K. Mong, F. Pollmann, Density matrix renormalization group on a cylinder in mixed real and momentum space. *Phys. Rev. B* **93**, 155,139 (2016). <https://doi.org/10.1103/PhysRevB.93.155139>
62. C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Natl. Bur. Stand. (U. S.)* **45**, 255–282 (1950)
63. E. Koch, *The Lanczos Method* (Forschungszentrum Jülich GmbH, 2011), chap. 8, pp. 8.1–8.30. <https://www.cond-mat.de/events/correl11/>
64. J. Haegeman, T.J. Osborne, F. Verstraete, Post-matrix product state methods: to tangent space and beyond. *Phys. Rev. B* **88**, 075,133 (2013). <https://doi.org/10.1103/PhysRevB.88.075133>
65. J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, F. Verstraete, Unifying time evolution and optimization with matrix product states. *Phys. Rev. B* **94**, 165,116 (2016). <https://doi.org/10.1103/PhysRevB.94.165116>
66. M. Fagotti, Entanglement & correlations in exactly solvable models. Ph.D. thesis, University of Pisa (2012). <https://etd.adm.unipi.it/t/etd-01312012-102006>
67. E. Lieb, T. Schultz, D. Mattis, Two soluble models of an antiferromagnetic chain. *Ann. Phys.* **16**(3), 407–466 (1961). [https://doi.org/10.1016/0003-4916\(61\)90115-4](https://doi.org/10.1016/0003-4916(61)90115-4)
68. M. Führinger, S. Rachel, R. Thomale, M. Greiter, P. Schmitteckert, DMRG studies of critical SU(N) spin chains. *Ann. Phys.* **520**(12), 922–936 (2008). <https://doi.org/10.1002/andp.20085201204>
69. Y. Ma, J. Wen, H. Ma, Density-matrix renormalization group algorithm with multi-level active space. *J. Chem. Phys.* **143**(3), 034,105 (2015). <https://doi.org/10.1063/1.4926833>
70. N. Kaushal, J. Herbrych, A. Nocera, G. Alvarez, A. Moreo, F.A. Reboredo, E. Dagotto, Density matrix

- renormalization group study of a three-orbital hubbard model with spin-orbit coupling in one dimension. *Phys. Rev. B* **96**, 155,111 (2017). <https://doi.org/10.1103/PhysRevB.96.155111>
71. S.R. White, Density matrix renormalization group algorithms with a single center site. *Phys. Rev. B* **72**, 180,403 (2005). <https://doi.org/10.1103/PhysRevB.72.180403>
72. C. Hubig, I.P. McCulloch, U. Schollwöck, F.A. Wolf, Strictly single-site DMRG algorithm with subspace expansion. *Phys. Rev. B* **91**, 155,403 (2015). <https://doi.org/10.1103/PhysRevB.91.155115>
73. A. Gleis, J.W. Li, J. von Delft, Controlled bond expansion for density matrix renormalization group ground state search at single-site costs. *Phys. Rev. Lett.* **130**, 246,402 (2023). <https://doi.org/10.1103/PhysRevLett.130.246402>
74. A.J. Daley, C. Kollath, U. Schollwöck, G. Vidal, Time-dependent density-matrix renormalization-group using adaptive effective Hilbert spaces. *J. Stat. Mech. Theory Exp.* **04**, P04,005 (2004). <https://doi.org/10.1088/1742-5468/2004/04/p04005>
75. G. Vidal, Efficient simulation of one-dimensional quantum many-body systems. *Phys. Rev. Lett.* **93**, 040,502 (2004). <https://doi.org/10.1103/PhysRevLett.93.040502>
76. S.R. White, A.E. Feiguin, Real-time evolution using the density matrix renormalization group. *Phys. Rev. Lett.* **93**, 076,401 (2004). <https://doi.org/10.1103/PhysRevLett.93.076401>
77. U. Schollwöck, S.R. White, Methods for time dependence in DMRG. *AIP Conf. Proc.* **816**(1), 155–185 (2006). <https://doi.org/10.1063/1.2178041>
78. S.R. White, Minimally entangled typical quantum states at finite temperature. *Phys. Rev. Lett.* **102**, 190,601 (2009). <https://doi.org/10.1103/PhysRevLett.102.190601>
79. E.M. Stoudenmire, S.R. White, Minimally entangled typical thermal state algorithms. *New J. Phys.* **12**(5), 055,026 (2010). <https://doi.org/10.1088/1367-2630/12/5/055026>
80. I.P. McCulloch, M. Gulácsi, The non-abelian density matrix renormalization group algorithm. *EPL (Europhys. Lett.)* **57**(6), 852–858 (2002). <https://doi.org/10.1209/epl/i2002-00393-0>
81. A. Weichselbaum, Non-abelian symmetries in tensor networks: a quantum symmetry space approach. *Ann. Phys.* **327**(12), 2972–3047 (2012). <https://doi.org/10.1016/j.aop.2012.07.009>
82. L. Hogben (ed.), *Handbook of Linear Algebra* (CRC Press, Boca Raton, 2006)
83. M.A. Nielsen, I.L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010)