

An introduction to infinite projected entangled-pair state methods for variational ground state simulations using automatic differentiation

Jan Naumann^{1,†,★}, Erik Lennart Weerda^{2,‡,★}, Matteo Rizzi^{2,3}, Jens Eisert^{1,4} and Philipp Schmoll^{1,◦},

1 Dahlem Center for Complex Quantum Systems and Institut für Theoretische Physik, Freie Universität Berlin, Arnimallee 14, 14195 Berlin, Germany

2 Institute for Theoretical Physics, University of Cologne, 50937 Köln, Germany

3 Forschungszentrum Jülich, Institute of Quantum Control, Peter Grünberg Institut (PGI-8), 52425 Jülich, Germany

4 Helmholtz-Zentrum Berlin für Materialien und Energie, Hahn-Meitner-Platz 1, 14109 Berlin, Germany

★ Both first authors have contributed equally.

† j.naumann@fu-berlin.de,

‡ weerda@thp.uni-koeln.de,

◦ philipp.schmoll@fu-berlin.de

Abstract

Tensor networks capture large classes of ground states of phases of quantum matter faithfully and efficiently. Their manipulation and contraction has remained a challenge over the years, however. For most of the history, ground state simulations of two-dimensional quantum lattice systems using (infinite) projected entangled pair states have relied on what is called a time-evolving block decimation. In recent years, multiple proposals for the variational optimization of the quantum state have been put forward, overcoming accuracy and convergence problems of previously known methods. The incorporation of automatic differentiation in tensor networks algorithms has ultimately enabled a new, flexible way for variational simulation of ground states and excited states. In this work we review the state-of-the-art of the variational iPEPS framework, providing a detailed introduction to automatic differentiation, a description of a general foundation into which various two-dimensional lattices can be conveniently incorporated, and demonstrative benchmarking results.

Copyright attribution to authors.

This work is a submission to SciPost Physics Lecture Notes.

License information to appear upon publication.

Publication information to appear upon publication.

Received Date

Accepted Date

Published Date

1	Contents	
2	1 Introduction	3
3	2 Variational iPEPS	4
4	2.1 iPEPS setup	5
5	2.2 CTMRG backbone	6
6	2.2.1 Absorption of iPEPS tensors	7
7	2.2.2 Calculation of projectors	9
8	2.2.3 Convergence and CTMRG fixed-points	11
9	2.3 Energy expectation values	12
10	2.4 Automatic differentiation	12
11	2.5 Calculation of the gradient at the CTMRG fixed-point	15
12	2.6 Optimization	16
13	2.7 Pitfalls and practical hints	18
14	2.7.1 Iterative SVD algorithm	18
15	2.7.2 Stability of the CTMRG routine	18
16	2.7.3 Prevention of local minima	19
17	2.7.4 Recycling of environments	19
18	2.7.5 Analysing iPEPS data at finite bond dimensions	19
19	2.7.6 Degenerate singular values	19
20	3 Extension to other lattices	20
21	3.1 Honeycomb lattice	20
22	3.2 Kagome lattice	21
23	3.3 Square-Kagome lattice	23
24	3.4 Triangular lattice	25
25	3.5 Comments about different structures	26
26	4 Benchmarks and discussions	27
27	4.1 Comments on lower bounds in variational principles	27
28	4.2 Honeycomb lattice	28
29	4.3 Kagome lattice	28
30	4.4 Square-Kagome lattice	30
31	4.5 Triangular lattice	31
32	4.6 Comments on excited states	32
33	4.7 Comments on fermionic systems	32
34	5 Conclusion and prospects	32
35	5.1 CO ₂ -emissions table	33
36	Appendix: Background on automatic differentiation	35
37	A Adjoint functions and variables	35
38	B Automatic differentiation for complex variables	36
39	C The implicit function theorem and its use at the CTMRG fixed-point	36
40	D Automatic differentiation in the language of differential geometry	36
41	References	40

1 Introduction

Tensor networks are at the basis of a wealth of methods that are able to efficiently capture systems with many degrees of freedom, primarily in the context of interacting quantum systems, but also in a wide range of other fields. They have a long history: The beginnings can be seen [1] as originating from work on transfer matrices [2] for two-dimensional classical Ising models and methods of corner transfer matrices again in the context of classical spin models [3]. In more recent times, the rise of tensor networks to describe interacting quantum many-body systems can be traced back to at least two strands of research. On the one hand, the now famous *density matrix renormalization group* (DMRG) approach [4, 5] can be regarded as a variational principle over *matrix product states* [6–8], a particularly common class of one-dimensional tensor network states. What are called *finitely-correlated states* [9] have later been understood as a Heisenberg picture variant of essentially the same family of states. These families of quantum states could further be interpreted as basically parametrizing gapped phases of matter in one spatial dimension. In a separate development, *tensor trains* became a useful tool in numerical mathematics [10]. These strands of research had been developing independently for quite a while before being unified in a common language of *tensor networks* (TN) as it stands now as a pillar of research on numerical and mathematical quantum many-body physics [11–15].

Two-dimensional tensor networks, now known as *projected entangled pair states* [16], again have a long history. The intuition why they provide a good ansatz class for describing ground states of gapped quantum many-body Hamiltonians [17, 18] – as well as other families of states – is the same as for matrix product states: Such states are expected to be part of what is called the “*physical corner*” of the Hilbert space. These states feature local entanglement compared to the degrees of entanglement unstructured states would exhibit. Ground states of gapped phases of matter are thought to satisfy *area laws for the entanglement entropy* [15]. Even though some of the rigorous underpinning of this mindset is less developed in two spatial dimensions compared to the situation in one spatial dimension, there is solid evidence that projected entangled pair states provide an extraordinarily good and powerful ansatz class for meaningful states of two-dimensional quantum systems.

There is a new challenge arising in such two-dimensional tensor networks. In contrast to matrix product states, they cannot be exactly efficiently *contracted*: On general grounds, there are complexity theoretic obstructions against the efficient contraction of projected entangled pair states in worst case [19] – and even in average case [20] – complexity. The burden can be lessened by acknowledging that projected entangled pair states can be contracted in quasi-polynomial time [21]. These more conceptual insights constitute an underpinning of a quite practically minded question: This shows that to develop ways of efficiently and feasibly approximating tensor network contractions in two spatial dimensions is at the heart of the method development in the field.

Consequently, over the years, several numerical methods of approximately contracting projected entangled pair states have been developed. In fact, much of the method development has been along these lines. In the focus of attention in this work are projected entangled pair states directly in the thermodynamic limit, commonly referred to as *infinite projected entangled pair states* (iPEPS) [22–24]. The contraction necessary to compute expectation values of local observables gives rise to the challenge of approximately calculating effective environments. Over the years, several methods have been introduced and pursued, including methods based on boundary matrix product states [22], corner transfer matrix methods [24–26] – particularly important for the method development presented here – and tensor coarse-graining techniques [27–30].

Variational optimization algorithms for uniform matrix product states have been developed

91 that combine density matrix renormalization group methods with matrix product state tangent
 92 space concepts to find ground states of one dimensional quantum lattices in the thermody-
 93 namic limit [31, 32], building on earlier steps of devising geometrically motivated variational
 94 principles for tensor network states [33, 34]. The pursuit of such variational optimization has
 95 been particularly fruitful in the two dimensional case of iPEPS. Initially proposed methods
 96 constructed the gradient of the energy explicitly using specialized environments [35, 36].

97 Recently, as an element of major method development, the programming technique called
 98 *automatic differentiation*, widely used in the machine learning community, has been utilized for
 99 the task of calculating the gradient [37] in tensor network optimization. This step drastically
 100 simplifies the programming involved and allows one to use variational ground state search on,
 101 e.g., more exotic lattice geometries with little additional effort. Such variational approaches
 102 for iPEPS constitute the basis for this work. Automatic differentiation has also been employed
 103 in further fashions in the tensor network context in several works recently [38–41, 41–49],
 104 some of which are accompanied by publicly available code libraries [50–53]. Notably, even for
 105 gapped local Hamiltonians with chiral topological ground states, for which the numerical appli-
 106 cability of PEPS was unclear due to no-go theorems in related cases [54], the use of variational
 107 optimization has proven successful [41, 49, 55]. As a novel programming paradigm, automatic
 108 differentiation composes parameterized algorithmic components in such a way that the pro-
 109 gram becomes differentiable and its components can be optimized using gradient search. It is
 110 a sophisticated way to evaluate the derivative of a function specified by a computer program,
 111 specifically by applying the chain rule to elementary arithmetic operations. Again, it has only
 112 recently been appreciated how extremely powerful such tools are in the study of interacting
 113 quantum matter by means of tensor networks.

114 In this review article, we elaborate on these developments and comprehensively present
 115 ideas for a variational iPEPS method based on automatic differentiation. This includes a de-
 116 tailed description of the methodology and practical insights for implementations, complement-
 117 ing and extending the existing body of literature. We further introduce a versatile framework,
 118 that allows arbitrary unit cells and different two-dimensional lattices to be treated on a com-
 119 mon footing. At the same time, this work accompanies the publicly available numerical library
 120 *variPEPS* – a versatile tensor network library for variational ground state simulations in two
 121 spatial dimensions – which implements the methods described in this review [56–58].

122 The content of this work is organised in three main sections. In Sec. 2, we describe the
 123 central methods that are being used in the variational iPEPS framework as well as practical
 124 remarks regarding implementation. Furthermore, we explain in detail the basics of automatic
 125 differentiation and its application in state-of-the-art ground-state search. In Sec. 3, we then
 126 turn to explaining how to conveniently map generic lattice structures to a square one, over
 127 which the variational iPEPS methods naturally operate. Following up on this, in Sec. 4, we
 128 present numerical benchmarks obtained with the methods outlined in the previous sections
 129 and implemented in the *variPEPS* library, in comparison to other customary methods like exact
 130 diagonalization, iPEPS imaginary-time evolution and variational Monte Carlo methods.

131 2 Variational iPEPS

132 We seek to find the the TN representation of the state vector $|\psi\rangle_{\text{TN}}$ that best approximates the
 133 true ground state vector $|\psi_0\rangle$ of an Hamilton operator of the form

$$H = \sum_{j \in \Lambda} T_j(h), \quad (1)$$

134 where T_j is the translation operator on the lattice Λ , and h is a generic k -local Hamiltonian,
 135 i.e., it includes an arbitrary number of operators acting on lattice sites at most at a (lattice)

136 distance k from a reference lattice point. Such a situation is very common in condensed matter
 137 physics, to say the least. To this aim, we employ the variational principle

$$\frac{\langle \psi | H | \psi \rangle}{\langle \psi | \psi \rangle} \geq E_0 \quad \forall |\psi\rangle, \quad (2)$$

138 and use an energy gradient with respect to the tensor coefficients to search for the minimum
 139 – the precise optimization strategy being discussed later. Such an energy gradient is accessed
 140 by means of tools from *automatic differentiation* (AD), a set of techniques to evaluate the
 141 derivative of a function specified by a computer program that will be summarized below. Since
 142 we directly target systems in the thermodynamic limit, a *corner transfer matrix renormalization*
 143 *group* (CTMRG) procedure constitutes the backbone of the algorithm, and also will come in
 144 handy for AD purposes. This is used to compute the approximate contraction of the infinite
 145 lattice, which is crucial in order to compute accurate expectation values in the first place.
 146 Importantly, the CTMRG routine is *always* performed on a regular square lattice, for which it
 147 can be conveniently defined. Support for other lattices, also non-bipartite ones, is possible by
 148 different lattice mappings, as we will demonstrate.

149 The method we will present in this section gives rise to an upper bound of the ground
 150 state energy in the sense of the variational principle as stated in Eq. (2). But we wish to point
 151 out at this point that for that to be strictly true it would be necessary to choose the CTMRG
 152 refinement parameter χ_E , introduced in detail in Sec. 2.2, to be $\chi_E \rightarrow \infty$. However, in practice
 153 we increase this refinement parameter χ_E until all observables are converged.

154 2.1 iPEPS setup

155 As introduced in the last section, we aim to simulate quantum many-body systems directly in
 156 the thermodynamic limit. To this end, we consider a unit cell of lattice sites that is repeated
 157 periodically over the infinite two-dimensional lattice. Reflecting this, the general configura-
 158 tions of the iPEPS ansatz are defined with an arbitrary unit cell of size (L_x, L_y) on the square
 159 lattice. The lattice setup, denoted by \mathcal{L} , can be specified by a single matrix, which uniquely
 160 determines the different lattice sites as well as their arrangement. Let us consider a concrete
 161 example of an $(L_x, L_y) = (2, 2)$ state with only two and all four individual tensors, denoted by

$$\mathcal{L}_1 = \begin{pmatrix} A & B \\ B & A \end{pmatrix}, \quad \mathcal{L}_2 = \begin{pmatrix} A & C \\ B & D \end{pmatrix}. \quad (3)$$

The corresponding iPEPS ansätze are visualized in Fig. 1. Here, the rows/columns of \mathcal{L} cor-




Figure 1: iPEPS ansätze with a unit cell of size $(L_x, L_y) = (2, 2)$ and only two (left) and four (right) different tensors as defined in Eq. (3).

162 respond to the x/y lattice directions. The unit cell \mathcal{L} is repeated periodically to generate the
 163 full two-dimensional system. As usual, the bulk bond dimension of the iPEPS tensors, denoted
 164 by χ_B , controls the accuracy of the ansatz. An iPEPS state with N different tensors in the unit
 165 cell consists of $Np\chi_B^4$ variational parameters, which we aim to optimize such that the iPEPS
 166




Figure 2: The norm of an iPEPS (here with a single-site unit cell) at a bulk bond dimension χ_B is approximated by a set of eight fixed-point environment tensors. The environment bond dimension χ_E controls the approximations in the CTMRG routine.

167 wave function represents an approximation of the ground state of a specific Hamiltonian. The
 168 parameter p denotes the dimension of the physical Hilbert space, e.g., $p = 2$ for a system of
 169 spin-1/2 particles.

170 The right choice of the unit cell is crucial in order to capture the structure of the targeted
 171 state. A mismatch of the ansatz could not only lead to a bad estimate of the ground state, but
 172 also to no convergence in the CTMRG routine at all. Different lattice configurations have to
 173 be evaluated for specific problems to find the correct pattern.

174 To circumvent the problem of a fixed and a priori chosen unit cell structure, recently an
 175 alternative description to the periodic structure has been proposed [59]. This approach is
 176 applicable if the Hamiltonian has a certain global symmetry, where the additional degree of
 177 freedom can be employed to reduce the description of the state to a subspace, e.g. $SU(2)$ for
 178 spin-1/2 systems. Here the state is described by the smallest possible unit cell, i.e. a single site
 179 for a square lattice, as well as a product of local unitary operators parameterized by a wave
 180 vector $\mathbf{k} = (k_x, k_y)$. A fixed choice of the wave vector then corresponds to the specification
 181 of a unit cell structure in the common iPEPS setup. This approach allows for a variational
 182 optimization of the wave vector along with the translationally invariant iPEPS tensor, removing
 183 the need to choose a fixed unit cell structure altogether.

184 In this work we restrict the description of the method to the common iPEPS setup with
 185 not only trivial unit cells. This enables the adaption of the framework to arbitrary, in general
 186 non-symmetric Hamiltonian models.

187 2.2 CTMRG backbone

188 One major drawback of two-dimensional TNs such as iPEPS is that the contraction of the
 189 full lattice can only be computed approximately. This is due to complexity theoretic obstruc-
 190 tions [19,20] and – practically speaking – the lack of a canonical form, which can only be found
 191 in loop-free tensor networks, for instance in matrix product states [8]. In order to circumvent
 192 the unfeasible exact contraction of the infinite 2d lattice, we employ an approximation scheme,
 193 the directional *corner transfer matrix renormalization group* (CTMRG) routine for iPEPS states
 194 with arbitrary unit cells of size (L_x, L_y) . The CTMRG method approximates the calculation
 195 of the norm $\langle \psi | \psi \rangle$ of the quantum state on the infinite square lattice by a set of effective
 196 environment tensors. This is achieved by an iterative coarse-graining procedure, in which all
 197 (local) iPEPS tensors in the unit cell \mathcal{L} are successively absorbed into the environment ten-
 198 sors towards all lattice directions, until the environment converges to a fixed-point. We will
 199 present a summary of the directional CTMRG methods for an arbitrary unit cell, following
 200 the state-of-the-art procedure [60–62]. The effective environment is displayed in Fig. 2, here
 201 for simplicity for a square lattice with a single-site unit cell $\mathcal{L} = (A)$. It consists of a set of
 202 eight fixed-point tensors, four corner tensors $\{C_1, C_2, C_3, C_4\}$ as well as four transfer tensors




Figure 3: Main steps of a left CTMRG move. One column of tensors is inserted into the network. Upon absorption of these tensors, the environment bond dimension grows rapidly, requiring a renormalisation step.

203 $\{T_1, T_2, T_3, T_4\}$, the latter sometimes also called edge tensors. In case of a larger unit cell, such
 204 a set of eight environment tensors is computed for each individually specified iPEPS tensor in
 205 the unit cell. The unavoidable approximations in the environment calculations are controlled
 206 by a second refinement parameter, the environment bond dimension χ_E .

207 In one full CTMRG step, the complete iPEPS unit cell is absorbed into the four lattice di-
 208 rections, such that the eight CTMRG tensors are updated for every iPEPS tensor. This is done
 209 column-by-column or row-by-row, depending on the direction. In each absorption step the
 210 environment bond dimension χ_E grows by a factor of χ_B^2 . To avoid an exponential increase
 211 in memory consumption and computation time, we need a method to truncate the bond di-
 212 mension back to χ_E . In order to do this, we calculate renormalization projectors for each row
 213 or column. Projectors are computed from a suitable patch of the iPEPS state including the
 214 effective environments, to find a best-possible truncation of the bond dimension. Different ap-
 215 proaches for their calculations have been proposed in the literature, which we will discuss in
 216 detail below, especially in the context of AD. In the following description of the CTMRG proce-
 217 dure we focus on a left absorption move, which grows all left environment tensors $\{C_4, T_4, C_1\}$.
 218 The main steps of insertion, absorption and renormalization are shown in Fig. 3. In Sec. 2.2.1,
 219 we will explain the full absorption procedure including renormalization, as it is done in prac-
 220 tise. Although projectors need to be calculated before the absorption, their motivation and the
 221 calculation of different projects is discussed later in Sec. 2.2.2.

222 2.2.1 Absorption of iPEPS tensors

223 In order to generate the CTMRG environment tensors, such that they converge to a fixed-
 224 point eventually, the iPEPS tensors are absorbed into them. To this end, we start with the
 225 network of one iPEPS tensor in the unit cell and its accompanying environment tensors. This
 226 is depicted in Fig. 3 in the top left. As shown on the top right of this figure, the network
 227 is extended by inserting one column, consisting of an iPEPS tensor and the top and bottom
 228 transfer tensors. While we depict the case of a single-site unit cell in Fig. 3, we note that
 229 the column of tensors to be inserted is generally dictated by the unit cell structure of the
 230 iPEPS ansatz, i.e., the left neighbor with the corresponding environment tensors for a left
 231 move. This crucial positional information for multi-site unit cells is specified by the coordinate

232 superscripts in the descriptions below. As indicated by the dashed line in Fig. 3, we absorb
 233 the inserted column into the left environment tensors by contracting all left pointing edges.
 234 This yields new environment tensors whose bond dimensions have grown by a factor χ_B^2 due
 235 to the virtual iPEPS indices, thus we need a way to truncate the dimension back to the CTMRG
 236 refinement parameter χ_E . This is done using the projectors we will discuss and compute in
 237 the next section. For now we introduce them as abstract objects labeled P that implement
 238 the dimensional reduction (i.e., the renormalization step) in an approximate but numerically
 feasible way. The updated tensor C'_1 is then given by the contraction in Fig. 4. As discussed




Figure 4: Update of the corner tensor C_1 in a left CTMRG step.

239 before, the correct tensors and projectors have to be used in accordance with the periodicity
 240 of the unit cell. The iPEPS tensor is now absorbed into the left transfer matrix T'_4 , where two
 241 projectors are needed to truncate the enlarged environment bond dimension. This is visualized
 242 in Fig. 5. Finally, the lower corner tensor C'_4 is updated, by absorbing a transfer matrix T_3 and




Figure 5: Update of the transfer matrix T_4 in a left CTMRG step. Here the projectors generally belong to different subspaces, unless the system is one-site translational invariant.

243 using another projector. The three absorption steps in Figs. 4, 5 and 6 are performed for all




Figure 6: Update of the corner tensor C_4 in a left CTM step.

244 rows x at a fixed column y , before moving to the next column $y + 1$. The process of computing
 245 projectors and growing the environment tensors is repeated for each column of the iPEPS unit cell,
 246 until the complete unit cell of $L_x \times L_y$ tensors has been absorbed into the left environment.
 247 This yields updated tensors C'_1 , T'_4 and C'_4 for all $[x, y]$.

249 The absorption of a full unit cell is then performed for the other three directions. In a top
 250 move the tensors C_1 , T_1 and C_2 are grown, in a right move the tensors C_2 , T_2 and C_3 and in
 251 a bottom move the tensors C_3 , T_3 and C_4 . This completes a *single* CTMRG step, which is then




Figure 7: Network of 2×2 iPEPS tensors and the corresponding CTMRG tensors, used as a starting point to compute the truncation projectors. For a left CTMRG step the top and bottom part is contracted into the matrices ρ_T and ρ_B with dimension $(\chi_E \chi_B^2) \times (\chi_E \chi_B^2)$. The red dashed line indicates the bonds that are renormalized back to a bond dimension χ_E .

252 repeated in the directional procedure until convergence is reached. In Sec. 2.2.3 we discuss
 253 appropriate convergence measures.

254 2.2.2 Calculation of projectors

255 In order to avoid an exponential increase of the bond dimension while growing the environ-
 256 nment tensors, projectors are introduced to keep the bond dimension at a maximal value of
 257 χ_E . Here, we will describe a common scheme to compute those projectors [61] and discuss
 258 some properties of their use in combination with AD [42]. The task of finding good projectors
 259 essentially comes down to finding a basis for the virtual space, whose bond dimension we aim
 260 to reduce, that can be used to distinguish between “more and less important” sub-spaces. This
 261 way, we can ideally reduce the dimension while keeping the most important sub-space. In
 262 what follows, we consider the lattice environment of the virtual space that we aim to truncate
 263 using the CTMRG environment tensors. To this end, we use a *singular value decomposition*
 264 (SVD) to identify the basis, in which the bond is optimally truncated such that we keep the
 265 most relevant information of this lattice environment. The lattice environment that we con-
 266 sider is shown in Fig. 7, where the red dotted line identifies the bonds that we aim to optimally
 267 truncate, illustrated for the example of a left absorption step. The arrangement of the tensors
 268 in the network of Fig. 7 follows the unit cell definition \mathcal{L} . For the trivial, single-site unit cell
 269 $\mathcal{L} = (A)$, all four iPEPS tensors are the same. We note that for a larger unit cell, cf. Fig. 1, the
 270 iPEPS tensors and their adjacent environments have to be chosen according to its periodicity.
 271 This setup for the arrangement is favorable, since it incorporates the (approximated) effect of
 272 the infinite environment by including all CTM tensors for the different lattice directions.

273 The projectors are used to renormalize the three left open tensor indices with combined
 274 bond dimension $\chi_E \chi_B^2$ back to the environment bond dimension χ_E in a left absorption step.
 275 In order to compute them, we start by defining the matrix

$$\mathcal{M} = \rho_B \cdot \rho_T \quad (4)$$

276 that represents the lattice environment of the virtual bond that we would like to truncate, as
 277 visualized in Fig. 8.

278 The procedure outlined here aims to find projectors P_{LT} and P_{LB} , such that the truncated
 279 matrix

$$\mathcal{M}_{\text{trunc}} = \rho_B \cdot P_{LT} \cdot P_{LB} \cdot \rho_T, \quad (5)$$




Figure 8: Matrix \mathcal{M} as defined by Eq. (4) in graphical TN notation. The red dashed line indicates the bonds that are renormalized back to a bond dimension χ_E .

280 is an optimal approximation to \mathcal{M} . To achieve this, we perform a singular value decomposition
281 on \mathcal{M} , i.e.,

$$\mathcal{M} = U_L S_L V_L^\dagger. \quad (6)$$

282 This factorization introduces a basis which allows for a separation of more relevant and less
283 relevant sub-spaces. To this end, we choose the largest χ_E singular values and their corre-
284 sponding singular vectors for the construction of the projectors. Furthermore, we define

$$S_L^+ = \text{inv}(\sqrt{S_L}), \quad (7)$$

285 where a pseudo-inverse with a certain tolerance is used. To increase the numerical stability, a
286 threshold of typically 10^{-6} (corresponding to a threshold of 10^{-12} for the singular values) is
287 used. Smaller singular values are set to zero. The use of a pseudo-inverse in the generation
288 of the projectors is equivalent to the construction of a projector with lower environment bond
289 dimension. Finally, the projectors to renormalize the left absorption step are construed as

$$\begin{aligned} P_{LT} &= \rho_T \cdot V_L \cdot S_L^+, \\ P_{LB} &= S_L^+ \cdot U_L^\dagger \cdot \rho_B. \end{aligned} \quad (8)$$

290 Here ρ_T and ρ_B again denote the top and bottom part of \mathcal{M} as introduced in Fig. 7. We
291 would like to point out the fact that without a truncation in the SVD above, the product of the
292 projectors we create in this way assembles the identity

$$\begin{aligned} P_{LT} \cdot P_{LB} &= \rho_T \cdot V_L \cdot S_L^{-1} \cdot U_L^\dagger \cdot \rho_B \\ &= \rho_T \cdot (\rho_B \cdot \rho_T)^{-1} \cdot \rho_B = \mathbb{1}. \end{aligned} \quad (9)$$

293 We stress again, that the choice of truncation in the calculations of the projectors is optimal in
294 order to approximate the lattice environment \mathcal{M} . A graphical representation of these projec-
295 tors is given in Fig. 9.

296 During a left-move, described in the previous section, we absorb the iPEPS tensors in the
297 unit cell column-by-column into the left environments. A renormalization step is required for
298 each of those moves, resulting in projectors that are specific to every bond. We therefore label
299 them by the positions in the unit cell, i.e., $P_{LT}^{[x,y]}$ and $P_{LB}^{[x,y]}$.

300 The process to generate the projectors described above uses the full lattice environment
301 \mathcal{M} , and thus we call them *full projectors*. It should be noted that Fishman et al. have proposed
302 a scheme to calculate equivalent projectors in a fashion that is numerically more stable, at the
303 cost of being computationally more expensive [62]. Their method is particularly useful in the
304 case of a singular value spectrum of \mathcal{M} that decays very fast.

305 Finally, different lattice environments of the virtual bond in question can be used to gener-
306 ate projectors. A very practical version are the so called *half projectors*. For those we choose a
307 lattice environment as illustrated in Fig. 10. These projectors are computationally less costly,
308 as they require a smaller network to be contracted. They only take into account correlation

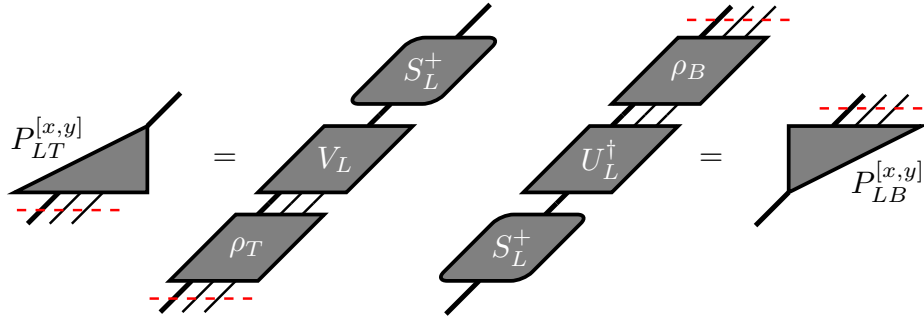


Figure 9: Calculation of top and bottom projectors for a left CTMRG absorption step. The red dashed line indicates the bonds that are renormalized back to a bond dimension χ_E .




Figure 10: Network of 2×1 iPEPS tensor and the corresponding CTMRG tensors, which is used as a reduced network to calculate the half projectors for a left CTMRG step. The red dashed line indicates the bonds that are renormalized back to a bond dimension χ_E .

309 within one half of the network, however this proves to be sufficient in many different applica-
 310 tions. Lately, there have been proposals for even cheaper alternatives of lattice environments
 311 and projector calculations [63], which yet have to be tested in the context of automatic differ-
 312 entiation and variational iPEPS optimization.

313 2.2.3 Convergence and CTMRG fixed-points

314 The CTMRG routine as described above is a power-method that eventually converges to a
 315 fixed-point. At this fixed-point, the set of environment tensors describes the contraction of
 316 the infinite lattice with an approximation controlled by the environment bond dimension χ_E .
 317 Convergence of the CTMRG tensors to the fixed-point can be monitored in different ways.
 318 In regular applications (those that do not involve automatic differentiation and gradients) the
 319 singular value spectrum of the corner tensors is typically a good quantity. Once the norm differ-
 320 ence of the spectrum between two successive CTM steps converges below a certain threshold,
 321 the environment tensors are assumed to be converged.

322 One peculiarity that is however not incorporated in this convergence check is sign or phase
 323 fluctuation for real or complex tensor entries, respectively. This means that, while projectors
 324 and hence the CTMRG tensors converge in absolute value, their entries can have different
 325 signs/phases in consecutive CTM steps. For reasons that become clear in Sec. 2.5 it is however
 326 required to reach *element-wise convergence* in the environment tensors for them to represent
 327 an actual fixed-point [42]. Those fluctuations originate from the gauge freedom in the SVD
 328 performed in Eq. (6). This is reflected in the freedom of introducing a unitary (block-)diagonal
 329 matrix Γ in an SVD,

$$\mathcal{M} = USV^\dagger = (U\Gamma)S(\Gamma^\dagger V^\dagger), \quad (10)$$

330 which leaves the expression invariant. The gauge freedom from the SVD directly affects the cal-
 331 culation of the projectors, such that we aim to fix the phases while computing these projectors.
 332 By eliminating this gauge freedom, at the true fixed-point, both projectors and environment
 333 tensors should be converged element-wise.

334 To fix the gauge, we introduce a diagonal unitary matrix Γ that redefines the phase of
 335 the largest entry (in absolute value) of every left singular vector to place it on the positive
 336 real axis [42]. To avoid instabilities of this gauge-fixing procedure due to numerical quasi-
 337 degeneracies, we always pick the first of such largest elements in basis order. Other choices,
 338 like addressing the first element with magnitude above a fixed threshold, are also possible.
 339 We further note that an alternative scheme to archive a fixed point in the CTMRG has recently
 340 been proposed [64].

341 2.3 Energy expectation values

342 Computing the energy expectation value required for the energy minimization is straightfor-
 343 ward using the CTMRG environment tensors. Assuming a Hamiltonian with only nearest-
 344 neighbour interaction terms, individual bond energies can be computed as shown in Fig. 11.
 345 The full energy expectation value, $\langle \psi | H | \psi \rangle / \langle \psi | \psi \rangle$, is obtained by collecting all different en-
 346 ergy contributions, i.e., all different terms in the Hamiltonian. Longer-range interaction can
 347 be treated as well, by simply enlarging the diagrams of Fig. 11 and performing more expensive
 348 contractions, which however occur only once per optimization step. In order to formulate a
 349 variational optimization of the tensor coefficients parametrizing the wave function, a gradient
 350 for the energy expectation value – including the foregone fixed-point CTMRG routine – is re-
 351 quired. This is achieved by the concept of automatic differentiation, as we will describe next.




Figure 11: Expectation values of a (horizontal) nearest-neighbour Hamiltonian term $\langle \psi | h_{i,j} | \psi \rangle / \langle \psi | \psi \rangle$ in tensor network notation, using the fixed-point CTMRG environments.

352

353 2.4 Automatic differentiation




Figure 12: Example of a computational graph for the function decomposition in Eq. (11).

354 *Automatic differentiation* (AD), sometimes also referred to as *algorithmic differentiation* or
 355 *automated differentiation*, is a method for taking the derivative of a complicated function which

is evaluated by some computer algorithm. It has been an important tool for optimization tasks in machine learning for many years. An introduction can be found in e.g. Ref. [65]. After its initial introduction in a foundational work [37], AD has found increasing applications in numerical TN algorithms in recent years [38, 39, 41, 42, 44, 45]. For the sake of simplicity, let us consider a function $E : \mathbb{R}^n \rightarrow \mathbb{R}^m$ for which we would like to evaluate the derivative. Noticeably, extensions to complex numbers are possible, and we provide some additional comments in Appendix B. We have the particular use-case of the energy expectation value $E(|\psi\rangle) = \langle \psi | H | \psi \rangle / \langle \psi | \psi \rangle$ of an iPEPS in mind, in which case the co-domain of the function E is \mathbb{R} . As we explain below, this has some important consequences for the use of AD.

Automatic differentiation makes use of the fact that many functions and algorithms are fundamentally built by concatenating elementary operations and functions like addition, multiplication, projection, exponentiation and taking powers, whose derivatives are known. The central insight is now that we can build up the gradient of a more complicated function from the derivatives of its elementary constituents by the *chain rule of differentiation*. In principle this even allows for a computation of the gradient to machine precision. It should be noted however, that it is neither necessary nor useful to deconstruct every function into its most elementary parts. Rather it is advantageous to deconstruct the function at hand only into a minimal amount of constituent-functions for which a derivative can be determined. These functions are often referred to as the *primitives* of the function of interest E . Primitives might themselves be a composition of many constituents but the derivative of the primitives themselves is known as a whole. An illustrative example for a primitive is a function that takes two matrices as an input and outputs the multiplication of them. On an elementary level this function is composed out of many multiplications and additions, but one can write down the derivative w.r.t. its inputs immediately. The choice of primitives describes the level of coarseness on which the AD process needs to know the details of the function E to compute the desired gradient. Defining large primitives of a function can reduce memory consumption, as well as increase performance and numerical stability of the AD process, e.g., by avoiding spurious divergences. Once the high-level function E has been decomposed into its minimal number of primitives, we can represent this decomposition with a so called *computational graph*. The computational graph is a directed, a-cyclic graph whose vertices represent the data generated as intermediate results by the primitives and the edges represent the primitives themselves, that transform the data from input to output.

As an example let us suppose we are able to decompose the function E into three primitives f_1, f_2 and f_3 , such that $E = f_3 \circ f_2 \circ f_1$. The primitives are maps between intermediate spaces

$$E : \mathbb{R}^{n_1} \xrightarrow{f_1} \mathbb{R}^{n_2} \xrightarrow{f_2} \mathbb{R}^{n_3} \xrightarrow{f_3} \mathbb{R}^{n_4} \quad (11)$$

and we refer to the variables in these spaces as $\vec{x}_i \in \mathbb{R}^{n_i}$. The computation graph illustrating this situation is shown in Fig. 12. AD can be performed in two distinct schemes, often called *forward-* and *backward-mode* AD. In the following we will demonstrate the two AD modes with the example of our previously introduced function E and its primitives. This will also serve to illustrate the computational cost of these AD schemes for the iPEPS use-case. Since f_1, f_2 and f_3 are said to be primitives, their Jacobians

$$J^i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_{i+1}} \times \mathbb{R}^{n_i},$$

$$J^i(\vec{x}_i^0) = \left(\frac{\partial f_i}{\partial \vec{x}_i} \right) \Big|_{\vec{x}_i = \vec{x}_i^0} \quad (12)$$

are known. An AD evaluation of the gradient of E at a specific point \vec{x}_1^0 is then given by the chain rule, the concatenation of the Jacobians of the primitives

$$\nabla E(\vec{x}_1^0) = J^3(\vec{x}_3^0) \cdot J^2(\vec{x}_2^0) \cdot J^1(\vec{x}_1^0), \quad (13)$$

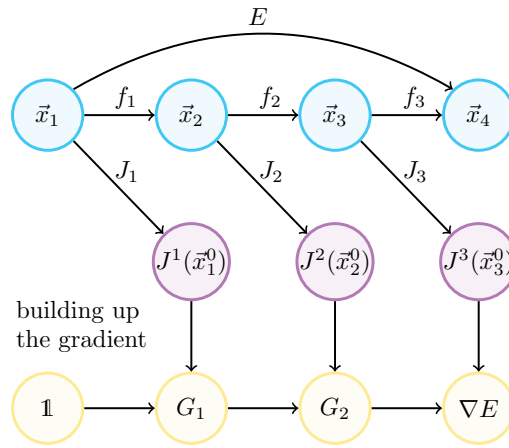


Figure 13: Illustration of forward-mode AD as described in Eq. (14) for the function decomposition in Eq. (11).

398 with $f_i(\vec{x}_i^0) = \vec{x}_{i+1}^0$. The difference between the *forward-* and *backward-mode AD* essentially
 399 comes down to the question from which side we perform the multiplication of the Jacobians
 400 above.

401 In the *forward-mode AD* scheme, the gradient is built up simultaneously with the evaluation
 402 of the primitives f_1, f_2 and f_3 , according to the following prescription for the i -th step:

$$\begin{aligned} f_i(\vec{x}_i^0) &= \vec{x}_{i+1}^0 \\ G_i &= J^i(\vec{x}_i^0) \cdot G_{i-1} \end{aligned} \quad (14)$$

403 with the starting condition $G_0 := \mathbb{1}_{n_1 \times n_1}$ and the final result $G_3 = \nabla E(\vec{x}_1^0) \in \mathbb{R}^{n_4 \times n_1}$. We see
 404 that in this case we build up Eq. (13) from right to left or “along the computational graph”
 405 as illustrated in Fig. 13. At first sight, such a procedure offers the potential advantage of not
 406 requiring to store intermediate results of the primitives in memory. However, if the dimension
 407 of the input (domain of E) is much larger than the dimension of the output (co-domain of
 408 E) – as it is the case in our use-case of iPEPS – this procedure becomes computationally very
 409 heavy. Indeed, saving and multiplying the large Jacobians in Eqs. (14) is often impractical.
 410 Thus, it is common to split up the starting condition $G_0 := \mathbb{1}_{n_1 \times n_1}$ into the n_1 canonical basis
 411 vectors $\{\vec{e}_i\}_{i=1, \dots, n_1}$. The procedure to generate the gradient from Eq. (14) is then repeated n_1
 412 times, each iteration generating a single component i . In this case, each step of the process of
 413 generating a component of the gradient is done by calculating a *Jacobian-vector product (JVP)*,
 414 so that only the resulting vector has to be stored. In order to create the full gradient in this
 415 way we need to repeat the procedure n_1 times, and the cost of calculating the full gradient
 416 scales as $\mathcal{O}(n_1) \times \mathcal{O}(E)$, where $\mathcal{O}(E)$ is the cost of evaluating E .

417 The *backward-mode AD* scheme works instead by first evaluating the function E and storing
 418 all intermediate results of the primitives along the way, and by then applying the iterative
 419 prescription

$$\bar{G}_i = \bar{G}_{i+1} \cdot J^i(\vec{x}_i^0) \quad (15)$$

420 with the starting condition $\bar{G}_4 = \mathbb{1}_{n_4 \times n_4}$ and the final result $\bar{G}_1 = \nabla E(\vec{x}_1^0) \in \mathbb{R}^{n_4 \times n_1}$. In the
 421 AD literature the objects \bar{G}_i are called adjoint variables and the functions that map the ad-
 422 joint variable on to each other, defined by Eq. (15), are called adjoint functions. We refer
 423 to Appendix A for more details on the adjoint functions and adjoint variables. In some parts
 424 of the literature the adjoint functions are also called pullbacks, which can be understood by
 425 looking at AD in language of differential geometry, cf. Appendix D. We see that in this case we
 426 build up Eq. (13) from left to right or as graphically illustrated in Fig. 14. This scheme has




Figure 14: Illustration of backward-mode AD as described in Eq. (15) for the function decomposition in Eq. (11).

427 the advantage of being computationally much cheaper if the output (co-domain) dimension is
 428 smaller than the input (domain) dimension – precisely the situation of our iPEPS setup, with
 429 $n_1 = Np\chi_B^4$ and $n_4 = 1$. We indeed only need to compute *vector-Jacobian products* (VJP) when
 430 evaluating the gradient, and, moreover, the full gradient is computed at once, instead of just
 431 a single element at a time as in the forward-mode AD scheme. This is why the cost of calcu-
 432 lating the gradient of the energy expectation value with *backwards-mode* AD is $\mathcal{O}(1) \times \mathcal{O}(E)$,
 433 which is superior to the cost of *forward-mode* AD. However, since we need to save all inter-
 434 mediate results of the primitives along the way in order to compute the gradient, the memory
 435 requirement for this scheme is in principle unbounded. Fortunately, the fixed-point condition
 436 for the iPEPS environments can be used to guarantee that the memory remains bounded in
 437 our calculations, as we illustrate in the following section.

438 2.5 Calculation of the gradient at the CTMRG fixed-point

439 Computationally, the CTMRG routine represents the bottleneck of the full iPEPS energy func-
 440 tion. It involves many expensive contractions and SVDs. Moreover, it requires an a priori
 441 unknown number of CTMRG iterations to reach convergence of the environment tensors. This
 442 would be especially disadvantageous for the gradient evaluation using plain-vanilla backward-
 443 mode AD, since this would require unrolling all the performed CTMRG iterations and paying
 444 a memory consumption linear in their number. However, this can be avoided by leveraging
 445 that fact that the CTMRG iteration eventually converges to a fixed point, and this is precisely
 446 the condition under which the energy evaluation is then performed. As soon this fixed point
 447 is reached, all CTMRG iterations are identical, i.e., reproducing the converged environment
 448 tensors. We can, in this situation, get away with only saving intermediate results from such a
 449 converged CTMRG iteration. This reduces the memory requirements by a factor of the num-
 450 ber of CTMRG iterations that we perform [37]. We stress here that, for this approach to work,
 451 we must make sure that the CTMRG procedure reaches an actual fixed point, meaning that
 452 all CTMRG environment tensors are converged element wise as discussed in Sec. 2.2.3. The
 453 fixed-point equation can be written as

$$e^*(A) = c(A, e^*(A)), \quad (16)$$

454 where the function c is one full CTMRG iteration, A are the iPEPS tensors which are constant
 455 during the CTMRG procedure and $e^*(A)$ represents the CTMRG environment tensors at the
 456 fixed-point. \mathcal{E} is the function that maps the iPEPS tensors with the fixed point environment




Figure 15: Computational graph of the CTMRG procedure for calculating the energy density at fixed point.

457 tensors and the Hamiltonian operators to the energy expectation value. The computational
 458 graph for the ground state energy is illustrated in Fig. 15. From it we can construct the form
 459 of the gradient of the energy expectation value with respect to the parameters of the iPEPS
 460 tensors A ,

$$\frac{\partial \langle H \rangle}{\partial A} = \frac{\partial \mathcal{E}}{\partial A} + \frac{\partial \mathcal{E}}{\partial e^*} \sum_{n=0}^{\infty} \left(\frac{\partial c}{\partial e^*} \right)^n \frac{\partial c}{\partial A}. \quad (17)$$

461 In practice this infinite sum is evaluated to finite order until the resulting gradient is converged
 462 to finite accuracy. An alternative viewpoint on the gradient at the fixed-point of the CTMRG
 463 procedure is presented in the Appendix C. It has recently been noted in Ref. [64] that the sta-
 464 bility and accuracy of the SVD derivative can be improved by including a previously neglected
 465 gradient contribution from the truncated part of the singular value spectrum.

466 2.6 Optimization

467 As discussed in the introduction of Sec. 2 we seek to find the iPEPS approximation $|\psi\rangle_{\text{TN}}$ of the
 468 ground state vector $|\psi_0\rangle$. Employing the methods discussed in the last sections we can describe
 469 this energy calculation as function $E(|\psi\rangle_{\text{TN}})$, consisting of the CTMRG power-method and the
 470 expectation value approximation using the resulting CTMRG environment tensors. Since we
 471 can calculate the gradient $\nabla E(|\psi\rangle_{\text{TN}})$ of this real scalar function it is straightforward to use
 472 well-known optimization methods to find the energy minimum. We would like to stress that
 473 the state vector $|\psi\rangle_{\text{TN}}$, and thus the energy function, only depends on the tensors defining
 474 the iPEPS ansatz and not the environment tensors since they are implicitly calculated from
 475 the ansatz. In this discussion we focus on two types of methods based on the gradient: The
 476 (*nonlinear*) *conjugate gradient* (CG) [66–70] and the quasi-Newton methods [71–76].

477 A naive approach to find the minimum of a function $E(|\psi_i\rangle)$, of which the gradient $\nabla E(|\psi_i\rangle)$
 478 is known, is to shift the input parameters $|\psi_i\rangle$ sufficiently along the negative gradient so that
 479 we find a new position $|\psi_{i+1}\rangle$ where the function value is reduced. At the end of this section
 480 we discuss what a sufficient step size means in this context. Iterating this procedure to a point
 481 where the gradient of the function vanishes (within a pre-defined tolerance) yields a solution

482 to the optimisation problem. Thus either a saddle point or a (local) minimum is reached then.
 483 This method is called steepest gradient descent. Although it resembles one of the simplest
 484 methods to find a descent direction, it is known to have a very slow convergence for difficult
 485 problems, e.g., for functions with narrow valleys [77]. Therefore, we use in practice more
 486 sophisticated methods to determine the descent direction.

487 The family of nonlinear conjugate gradient as generalization of the linear conjugate gra-
 488 dient method modifies this approach. Instead of using the negative gradient as a direction in
 489 each iteration step it uses a descent direction which is conjugated to the previous ones. For
 490 the linear conjugate gradient method there is a known factor β_i to calculate the new descent
 491 direction $d_i = g_i + \beta_i d_{i-1}$ from the gradient g_i of the current step and the descent direction
 492 d_{i-1} of last step. In the generalization for nonlinear functions this parameter is not uniquely
 493 determined anymore, however there are different approaches to estimate this parameter in the
 494 literature [67–69]. In our implementation we chose the nonlinear conjugate gradient method
 495 in the formulation as has been suggested by Hager and Zhang [70],

$$\begin{aligned}\tilde{\beta}_i^{\text{HZ}} &= \frac{1}{d_{i-1}^\top y_i} \left(y_i - 2d_{i-1} \frac{\|y_i\|^2}{d_{i-1}^\top y_i} \right)^\top g_i, \\ \eta_i &= \frac{-1}{\|d_{i-1}\| \min(\eta, \|g_{i-1}\|)}, \\ \beta_i^{\text{HZ}} &= \max(\tilde{\beta}_i^{\text{HZ}}, \eta_i),\end{aligned}\tag{18}$$

496 with $\|\cdot\|$ the Euclidian norm, $y_i = g_i - g_{i-1}$ and $\eta > 0$ a numerical control parameter which
 497 has been set to $\eta = 0.01$ in the work by Hager and Zhang. In our tests and benchmarks this
 498 choice for β_i has been proven to be numerically stable.

499 The other family of optimization methods we use in our implementation are the quasi-
 500 Newton methods, concretely the *Broyden–Fletcher–Goldfarb–Shanno* (BFGS) algorithm [73–
 501 76] and its *low-memory* (L-BFGS) variant [72, 78]. These methods are based on the Newton
 502 method where the descent direction is calculated using not only the gradient, but also the
 503 second derivative (the Hessian matrix). Unfortunately, it is computationally expensive to cal-
 504 culate the Hessian for large sets of input parameter, which makes this method only feasible
 505 for small parameter sets (i.e., iPEPS ansätze with a small number of variational parameters).
 506 Quasi-Newton methods solve this problem by not calculating the full Hessian, but an approx-
 507 imation of it. To this end, the gradient information from successive iteration steps is used to
 508 update the approximation in each step. The BFGS algorithm stores the full approximated Hes-
 509 sian matrix, including the information from all previous steps. In contrast, the L-BFGS method
 510 calculates the effective descent direction in an iterative manner from the last N optimization
 511 steps. This way not the full (approximated) Hessian has to be stored in memory but only the
 512 gradients of the last N steps. This reduces the memory consumption by an order of magnitude.
 513 The disadvantage is that not the full information of all previous steps is considered, but only
 514 a fraction of it. Nevertheless, due to the memory requirements to store the full approximated
 515 Hessian in the standard BFGS method for larger iPEPS bond dimensions we use L-BFGS as the
 516 default quasi-Newton method.

517 As noted before, we would like to shift the variational parameters x_i along the descent
 518 direction d_i determined by the different algorithms discussed above. With this shift we aim to
 519 find a new ansatz $x_{i+1} = x_i + \alpha_i d_i$ with α_i the step size along the descent direction. Ideally,
 520 we would like to find the optimal step size $\alpha_i = \min_\alpha E(x_i + \alpha d_i)$ minimizing the function
 521 value along the descent direction. However, determining this optimal value is computationally
 522 expensive and thus in practice, we stick to a sufficient step size fulfilling some conditions. The
 523 procedure to find this step size is called *line search* [79–82]. In our implementation we use
 524 the Wolfe conditions [80–82], since they guarantee properties which are feasible particularly
 525 for the (L-)BFGS method and its iterative update of the effect of the approximate Hessian.

526 2.7 Pitfalls and practical hints

527 2.7.1 Iterative SVD algorithm

528 We also advertise the use of iterative algorithms for the calculation of the SVD in the CTMRG
 529 procedure. This can be quite advantageous computationally, since only χ_E singular values are
 530 needed for a matrix of size $(\chi_E \chi_B^2) \times (\chi_E \chi_B^2)$ during the CTMRG. To this end we use the
 531 the *Golub-Kahan-Lanczos* (GKL) bidiagonalization algorithm with additional orthogonalization
 532 for the Krylov vectors. This algorithm is available, e.g., in packages like KRYLOVKIT.JL [83] or
 533 ITERATIVESOLVERS.JL [84] in the JULIA programming language. We highlight the utility of this
 534 type of algorithm for the calculation of the SVD with the comparison of the computational
 time of the different algorithms in the iPEPS use case in Fig. 16.




Figure 16: Comparison of the computational time for the calculation of the first χ_E singular values/vectors of a matrix of dimension $(\chi_E \chi_B^2) \times (\chi_E \chi_B^2)$ obtained in a CTMRG procedure with bond dimension $\chi_B = 6$. The conventional SVD (blue), which is truncated only after calculating the full SVD spectrum is substantially slower than the iterative GKL methods. The GKL algorithm in the CTMRG use case was showed comparable performance when constructing the $\chi_E \chi_B^2$ matrix explicitly (orange) or by just implementing its action of a vector (green). While the GKL algorithm for the case at moderate d and χ_E constructing the matrix usually is faster, at larger χ_B and χ_E it can become advantages to only implement the action of the matrix.

535

536 2.7.2 Stability of the CTMRG routine

537 One of the basic prerequisite for a stable variational iPEPS optimization is a robust CTMRG
 538 routine fulfilling the convergence requirements discussed in Sec. 2.2.3. Obviously, there is
 539 the environment bond dimension χ_E to control the accuracy of the approximation of the envi-
 540 ronment. If the environment bond dimension is chosen too low, the approximation is invalid
 541 and the CTMRG routine can yield an inaccurate result for the expectation value. This could
 542 further lead to an unstable variational update. To check heuristically whether the refinement
 543 parameter χ_E is chosen sufficiently high, one can check the singular value spectrum obtained
 544 during the projector calculation as described in Sec. 2.2.2. As a reliable criteria for the amount
 545 of information loss, we compute the truncation error ϵ_T given by the norm of the discarded

546 singular values of the normalized spectrum [85]. If the truncation error is larger than some
547 threshold (e.g., $\varepsilon_T > 10^{-5}$), one can assume that the environment bond dimension is chosen
548 too low and has to be increased. Employing this procedure, the bond dimension can automat-
549 ically be increased during the variational optimization if necessary. A sufficiently large χ_E is
550 crucial as the AD optimization can otherwise exploit the inaccuracies of the CTMRG procedure,
551 leading to false ground states with artificially low energy.

552 2.7.3 Prevention of local minima

553 An ideal iPEPS optimization finds the global energy minimum of the input Hamiltonian within
554 the iPEPS ansatz class of fixed unit cell and bond dimension. In practice, however, it is possible
555 – and likely – that the algorithm gets stuck in local minima. In order to avoid local minima
556 and reach the global optimum, there are a number of possible tricks. The naive way is to start
557 several simulations with different random initial states. This is typically a practicable solution,
558 although it is not well controllable and requires large computational resources.

559 An optimization of a system with a tendency for local minima might still be successful, if a
560 suitable initial state is provided. One possibility are initial states obtained by imaginary-time
561 evolution methods (simple update, full update [22,23,86]). While this is typically a convenient
562 solution, it is sometimes necessary to perturb the input tensors with a small amount of noise
563 (e.g., 10^{-2} in relative amplitude) to actually avoid local minima. As an alternative, one can
564 input a converged state obtained from energy minimization of a different TN ansatz, provided
565 there is a suitable mapping between the different structures. Examples for this technique are
566 provided for benchmarks on different lattices in Sec. 4.

567 Finally, the method of perturbing a suitable initial state with small amount of random
568 noise of course could also be applied to the result of one optimization run. As suggested in
569 the literature [87], this could help to escape possible local minima. Therefore, one could retry
570 this method a few times and keep the best result of all runs.

571 2.7.4 Recycling of environments

572 The calculation of the environment tensors with the CTMRG routine is expensive and time
573 consuming. During an optimization process one can reuse the environment tensors of the
574 previous optimization step as input for the next. This is advisable in the advanced stages of
575 the optimization, in which the gradient is already small. In this scenario the iPEPS tensors
576 usually only change minutely, such that starting the CTMRG routine from the environments
577 of the last PEPS tensor can reduce the number of CTMRG steps required for convergence
578 substantially.

579 2.7.5 Analysing iPEPS data at finite bond dimensions

580 Data generated with the variational iPEPS setup inevitably carries finite iPEPS bond dimension
581 χ_B (or even finite environment bond dimension χ_E) effects. Several schemes are available to
582 utilize the correlation length of the optimal tensors at a certain value of χ_B to extrapolate the
583 values of observables [88–90]. Additionally, a extrapolation scheme using data of an optimized
584 iPEPS state at finite χ_B and finite but suboptimal χ_E has been proposed and shown useful [91].

585 2.7.6 Degenerate singular values

586 Although very rare, a degenerate singular value spectrum in the calculation of the projec-
587 tors can be an obstacle. The gradient of the SVD becomes ill-defined in this case, due to terms
588 $F_{i,j} = 1/(s_j^2 - s_i^2)$ in the derivative [45], where s_i are the singular values. Naturally, it would be




Figure 17: Honeycomb and topologically equivalent brick-wall lattice.

589 desirable to remove the degeneracy by constraining the system to the correct physical symme-
 590 try, thereby grouping the degenerate singular values to common multiplets of the underlying
 591 symmetry group. If this is not possible or the degeneracies appear independently of a symme-
 592 try (“accidental” degeneracy), workarounds have to be used. One possibility is to add a small
 593 amount of noise in the form of a diagonal matrix XX^{-1} on the CTMRG environment links, with
 594 the elements of X drawn from a tiny interval $[1 - \varepsilon, 1 + \varepsilon]$. This can space out the singular
 595 value spectrum and stabilize the SVD derivative [92]. Recently an alternative procedure to
 596 eliminate divergences in the derivative of the SVD with degenerate spectrum has been pro-
 597 posed in Ref. [64]. Here, for the case of a rotationally invariant CTMRG, the divergent term
 598 is canceled out by a particular gauge fixing of the environment tensors.

599 3 Extension to other lattices

600 The directional CTMRG routine on the square lattice is very convenient for its orthogonal
 601 lattice vectors and definition of the effective environments. It is therefore natural to exploit
 602 the implemented routines for different kind of lattices that can be mapped back to the square
 603 lattice. This can typically be achieved by a suitable coarse-graining, in which a collection of
 604 lattice sites on the original lattice is mapped into an effective site on the square lattice. Energy
 605 expectation values can then be directly evaluated in the coarse-grained picture as well. This is
 606 even advantageous for the AD optimization procedure, since the energy can often be computed
 607 with a smaller number of individual terms. In this section we will present the mapping for four
 608 types of lattices frequently found in condensed matter systems – the honeycomb, Kagome,
 609 square-Kagome and triangular lattice. Naturally, the framework can be extended by other
 610 suitable two-dimensional lattices, such as dice, square-octagon, maple-leaf and others. As
 611 an alternative to the coarse-graining approach, CTMRG methods that directly operate on the
 612 original lattice structures can also be defined [46, 93].

613 3.1 Honeycomb lattice

614 The honeycomb, hexagonal or brick-wall lattice is of broad interest in material science and of-
 615 ten appears in the context of quantum many-body systems. For instance, the *Kitaev honeycomb*
 616 *model* is a paradigmatic example hosting different kinds of phases supporting different types
 617 of anyons, both Abelian and non-Abelian [94]. We will now describe the general technical
 618 framework to simulate honeycomb lattices with the backbone CTMRG procedure described
 619 in Sec. 2.2. To this end we consider an elementary unit cell of the honeycomb lattice. Here
 620 we choose to define it along so-called *x*-links for reasons that become clear soon. Alterna-
 621 tively and equivalently, it could as well be defined along *y*- or *z*-links. As an example with
 622 eight different tensors on the honeycomb lattice, corresponding to four elementary unit cells




Figure 18: iPEPS ansatz on the honeycomb lattice with four elementary unit cells, resulting in eight different lattice sites. x -, y - and z -links denote the three types of inequivalent links in the lattice. Coarse-graining this state to a square lattice results in a $(L_x, L_y) = (2, 2)$ configuration, with an arrangement as in Eq. (3) / Fig. 1.

623 is shown in Fig. 18. Coarse-graining the two lattice sites along x -links of the honeycomb lat-
 624 tice directly results in a square lattice, as shown in Fig. 19. Here, the (mapped) unit cell has
 625 size $(L_x, L_y) = (2, 2)$ with an arrangement as in Eq. (3) and Fig. 1. The green color is used
 626 to highlight the coarse-graining along x -links. In contrast to the regular square lattice, each
 627 coarse-grained tensor has two physical indices that can be reshaped to a single, combined index
 628 before feeding it into the CTMRG procedure. A trivial unit cell on the square lattice, consisting
 629 of only a single-site tensor, results in two different tensors on the honeycomb lattice.

630 The CTMRG routine can then be run as described above, just with a larger physical di-
 631 mension. This does not change anything in the contractions, it is just computationally more
 632 expensive. Expectation values can now be evaluated accurately using the CTMRG environ-
 633 ment tensors. Assuming nearest-neighbour terms again, expectation values along x -links can
 634 be computed by a single-site TN, while y - and z -bonds remain two-site TNs similarly to Fig. 11.

635 3.2 Kagome lattice

636 Another important and often encountered lattice in condensed matter physics is the Kagome
 637 lattice. It is of special interest due to its corner-sharing triangles, which lead a strong geomet-
 638 ric frustration for anti-ferromagnetic models. Using a simple mapping of the Kagome lattice
 639 to a square lattice, we can directly incorporate it into our variational PEPS library. The Kagome
 640 lattice is shown in Fig. 20a. Naturally, we can define a unit cell of tensors that is repeated
 641 periodically over the whole two-dimensional lattice. In our setting we consider an upward tri-
 642 angle on the Kagome lattice as an elementary unit cell, highlighted by the gray dotted area in
 643 Fig. 20a. By choosing a coarse-graining, we can represent the three lattice sites in the unit cell
 644 by a single iPEPS tensor, which connects to its neighbours by four virtual indices. This direct
 645 mapping is shown in Fig. 20b. Nearest-neighbour links in the Kagome lattice get mapped to
 646 nearest-neighbour or second-nearest-neighbour links in the square lattice. Every iPEPS site on




Figure 19: Using a mapping the brick-wall lattice is transformed to the square lattice. The green color of the tensors is just to highlight the coarse-graining along x -links, while y - and z - links remain in the network.




Figure 20: (a) Regular Kagome lattice with corner-sharing triangles and an elementary unit cell consisting of three lattice sites. (b) Regular Kagome lattice mapped to a square lattice by coarse-graining of the three spins in each unit cell.




Figure 21: Honeycomb lattice (dual to the Kagome lattice) with spins residing on the lattice links and additional simplex tensors on the lattice sites. Unit cells are highlighted by the gray dotted areas. Upon coarse-graining of the unit cells, the dual honeycomb lattice is mapped to the regular square lattice. Physical indices of the corresponding TN states are not shown.

647 the square lattice has a physical dimension of p^3 . As an alternative mapping, which results
 648 in the same coarse-grained TN structure, we move from the Kagome lattice to its dual, the
 649 honeycomb lattice. Here the spins live on the links instead of the vertices. The honeycomb
 650 mapping presented in Sec. 3.1 is therefore not directly applicable and additional simplex ten-
 651 sors are necessary to connect the lattices sites. This TN structure is shown in Fig. 21, which
 652 is commonly known as the infinite *projected entangled simplex state* (iPESS) [95]. Due to this
 653 particular mapping, three Kagome lattice sites (along with two simplex tensors) are coarse-
 654 grained into a single iPEPS site on the square lattice. While the mappings in Fig. 20b and
 655 Fig. 21 result in the same square lattice TN, they differ in the number of variational param-
 656 eters in the ansatz. In the direct iPEPS ansatz, every unit cell tensor has $p^3 \chi_B^4$ parameters,
 657 while there are only $(3p \chi_B^2 + 2 \chi_B^3)$ parameters for the iPESS ansatz. Moreover, quantum cor-
 658 relations between lattice sites are exactly captured within the coarse-grained cluster for the
 659 iPEPS, whereas they are limited by the bulk bond dimension for the iPESS. In the ladder case,
 660 however, there is no bias between lattice sites within one cluster and sites belonging to differ-
 661 ent clusters. The nearest neighbor interactions on the Kagome lattice are mapped to on-site,
 662 nearest neighbor and next-nearest neighbor interactions on the square lattice. As a concrete
 663 mapping example which has particular use in the study of the regular Heisenberg model in a




Figure 22: Kagome lattice structure corresponding to a square lattice unit cell according to Eq. (19).




Figure 23: Square-Kagome lattice. Similarly to the regular Kagome lattice, it features corner-sharing triangles. The elementary unit cell consists of six sites, as shown in Fig. 24.

664 magnetic field, we consider the iPEPS configuration

$$\mathcal{L} = \begin{pmatrix} A & B & C \\ B & C & A \\ C & A & B \end{pmatrix} \quad (19)$$

665 on the square lattice. This configuration results in the Kagome lattice structure shown in
 666 Fig. 22.

667 3.3 Square-Kagome lattice

668 As a third lattice that has gained a lot of interest in recent time is the square-Kagome lattice.
 669 Similar to the regular Kagome lattice it features corner-sharing triangles and it is expected to
 670 host exotic quantum phases due to the geometric frustration for antiferromagnetic spin models.
 671 The square-Kagome lattice structure is shown in Fig. 23. Naturally, a coarse-graining of the six
 672 spins in the elementary unit cell can be used, which directly maps the square-Kagome lattice
 673 to a square lattice as depicted in Fig. 24. Following the same construction as for the regular
 674 Kagome lattice, we can generalize the iPESS ansatz to the dual of the square-Kagome lattice,




Figure 24: Regular square-Kagome lattice mapped to a square lattice by coarse-graining the six spins in each elementary unit cell.




Figure 25: Square-octagon lattice (dual to the square-Kagome lattice) with spins residing on the lattice links and additional simplex tensors on the lattice sites. Unit cells are highlighted by the gray dotted areas. Upon coarse-graining of the unit cells, the square-octagon lattice is mapped to the regular square lattice. Physical indices of the corresponding TN states are not shown.

675 the so-called $(4, 8^2)$ Archimedean lattice. This results in an ansatz with four simplex tensors
 676 and six lattice site tensors per elementary unit cell, as illustrated in Fig. 25. Counting the
 677 number of variational parameters in both TN ansätze, we find a drastic reduction in the iPES
 678 ansatz, again. Here the iPEPS has $p^6 \chi_B^4$ parameters, while the iPES only has $(6p\chi_B^2 + 4\chi_B^3)$
 679 parameters for each tensor in the unit cell. In Table 1, we reinforce the difference for usual
 680 iPEPS bond dimensions, which has a strong influence on the expressivity and optimization
 681 of the different TN structures. As in the case of the Kagome lattice, the first coarse-graining
 682 captures quantum correlations within the cluster exactly. While this is not the case for the
 683 iPES mapping, it does not introduce a bias for the different lattice sites within and across
 684 clusters. Both mappings result in a large physical bond dimension of p^6 , with p the Hilbert
 685 space dimension of the original degrees of freedom (e.g., $p = 2$ for a spin-1/2). This makes
 686 especially the CTMRG routine computationally expensive. As an example we consider a two-
 687 site checkerboard pattern $((L_x, L_y) = (2, 2)$ with only two different tensors) on the square

χ_B	$p^6 \chi_B^4$	$(6p\chi_B^2 + 4\chi_B^3)$	ratio
2	1024	80	12.8
3	5184	216	25.0
4	16384	448	36.6
5	40000	800	50.0
6	82944	1296	64.0
7	153664	1960	78.4
8	262144	2816	93.1

Table 1: Number of variational parameters (per elementary unit cell) in the iPEPS and iPESSTN ansatz of the square-Kagome lattice for $p = 2$, assuming real tensor elements.




Figure 26: Square-Kagome lattice structure for a square lattice unit cell according to Eq. (20). The ansatz has twelve different lattice sites with two-site translation invariance in both x - and y -direction.

688 lattice, given by

$$\mathcal{L} = \begin{pmatrix} A & B \\ B & A \end{pmatrix}. \quad (20)$$

689 This results in a square-Kagome state with twelve different lattice sites, as shown in Fig. 26.

690 Assuming nearest-neighbour interactions in the Hamiltonian, the ground state energy can
691 be computed by single-site as well as horizontal and vertical two-site expectation values.

692 3.4 Triangular lattice

693 The triangular lattice, shown in Fig. 27 is another two-dimensional lattice variant that appears
694 frequently in condensed matter systems. Due to its large connectivity to six nearest neighbours,
695 it is a typical playground for frustrated systems, hosting a variety of different quantum phases.
696 As a consequence of this, the large connectivity makes it more challenging for numerical sim-
697 ulations. The triangular lattice can be directly interpreted as a square lattice with additional
698 diagonal interactions. The entanglement between diagonal sites is then mediated by the reg-
699 ular virtual links in the square lattice tensor network. Nearest-neighbour interactions on the
700 triangular lattice are again mapped to nearest-neighbour and next-to-nearest-neighbour inter-
701 action on the coarse-grained square lattice.




Figure 27: Regular triangular lattice with a connectivity of six, i.e., every lattice site is connected to six nearest neighbours.




Figure 28: iPESS ansatz for the triangular lattice consisting of only two tensors per triangular lattice site. When one lattice site and one simplex tensor are combined, the triangular lattice is directly mapped onto a regular square lattice.

702 An alternative TN representation of the triangular lattice can be constructed using again
 703 the iPESS ansatz. In contrast to the iPESS for Kagome and square-Kagome lattices, here the
 704 lattice sites have three virtual indices, too. The mapping is visualized in Fig. 28 with the iPESS
 705 ansatz being a honeycomb lattice. Similarly to the first interpretation, this iPESS honeycomb
 706 ansatz can be mapped to a regular square lattice with additional next-to-nearest-neighbour
 707 interactions. While the first approach as $p\chi_B^4$ parameters per unit cell tensor, the iPESS map-
 708 ping only has $(p\chi_B^3 + \chi_B^3)$ coefficients. Finally, and as an alternative to the previous mappings,
 709 a reverse transformation could be used, which involves a fine-graining of the lattice sites [96].

710 3.5 Comments about different structures

711 In general there is no unique way to map a given lattice structure to the square lattice. The
 712 different approaches mainly differ in the number of variational parameters. While the energy
 713 for an ansatz with fewer parameters can be optimized with fewer resources, an ansatz with a
 714 higher variational freedom might be able to capture the physical system more accurately. At
 715 the same time the optimization becomes more complex due to the need to calculate bigger gra-
 716 dients. In practice, choosing the right ansatz depends on the spatial structures of the quantum
 717 state, the amount of entanglement present in the system and the required accuracy. One strat-
 718 egy that works well is a step-wise optimization. In the first step one can choose, e.g., an iPESS
 719 ansatz with fewer variational parameters. Once an optimized wave function has been found,
 720 the iPESS ansatz is coarse-grained into a TN with a higher number of variational parameters,
 721 e.g., a direct iPEPS ansatz. A second optimization of this more expressive ansatz might then
 722 result in lower ground state energies. In the following sections we will present benchmarks,
 723 where several of the lowest data points have been obtained with such a two-step procedure.

724 4 Benchmarks and discussions

725 In this section, we will present benchmarks for a challenging and paradigmatic models on
 726 the different currently supported lattices. Due to its prominence and availability of bench-
 727 marks to different numerical techniques, we generally focus on the spin-1/2 Heisenberg anti-
 728 ferromagnet. The Heisenberg Hamiltonian is given by

$$H = J \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j, \quad (21)$$

729 where $\langle i, j \rangle$ denotes nearest neighbours and \vec{S}_i are the spin-1/2 operators on the lattice sites.
 730 We consider isotropic anti-ferromagnetic interactions at $J = 1.0$ throughout the benchmark
 731 section. Variational energies obtained with our implementation are denoted by “*variational*
 732 *update*” (VU). Where applicable, we include different TN variants (e.g., iPESS and iPEPS) in
 733 the numerical benchmarks, to highlight the effect of different numbers of variational param-
 734 eters. Imaginary time-evolution in the form of a “*simple update*” (SU) on the different lattice
 735 structures can provide initial states for the variational update as discussed in Sec. 2.7.3. When-
 736 ever we use initial tensors from the SU, we add a small amount of random noise to the input
 737 tensors prior to the variational update, in order to circumvent possible local minima in the
 738 imaginary time evolution.

739 In the plots of this section we include the energies calculated by the mean-field environ-
 740 ment (MF) used in the simple update. Using this approximation much larger iPEPS bond
 741 dimensions are computationally feasible but we would like to point out that this method is
 742 not guaranteed to be variational in the sense that the energy is an upper bound to the ground
 743 state energy. Thus, it is only sensible to rigorously compare results for which energy expecta-
 744 tion values are computed by CTMRG. We include the non-variational MF energies for higher
 745 iPEPS bond dimensions for a rough comparison.

746 We add for each benchmark a table with the comparison of the results obtained by the
 747 simple update simulations and the best result throughout all variational updates for a fixed
 748 iPEPS bond dimension χ_B . Both expectation values have been calculated by CTMRG.

749 4.1 Comments on lower bounds in variational principles


750 As a further conceptual point, it is important to stress that variational principles can be bench-
 751 marked as well by resorting to lower bounds to ground state energies. Such lower bounds can
 752 be efficiently computed and hold in the thermodynamic limit up to a small constant error in
 753 the energy density [97]. If the Hamiltonian H is seen as being written as a sum of terms

$$H = \sum_j h_j \quad (22)$$

754 where each h_j is a patch that contains as many unit cells that can be accommodated in an
 755 exact diagonalization, then

$$\frac{\langle \psi | H | \psi \rangle}{\langle \psi | \psi \rangle} \geq E_0 \quad \forall |\psi\rangle, \quad E_0 \geq \lambda_{\min}(h_j), \quad (23)$$

756 where $\lambda_{\min}(h_j)$ denotes the smallest eigenvalue of the patch h_j with open boundary conditions.
 757 In this way, the quality of the variational principle giving rise to upper bounds to the ground
 758 state energy can be certified by lower bounds.



χ_B	E_0 (SU)	E_0 (VU)
2	-0.53533	-0.537600
3	-0.53969	-0.541145
4	-0.54346	-0.544159
5	-0.54398	-0.544474
6	-0.54409	-0.544536
7	-0.54412	-0.544543


Figure 29: Benchmarking results for the isotropic spin-1/2 Heisenberg model on the honeycomb lattice. For comparison we include the variational result obtained by an iPEPS study in Ref. [98]. Additionally, the result calculated by the coupled cluster method in Ref. [99] is shown, which is due to extrapolation not variational either.

759 4.2 Honeycomb lattice

760 For the simulations of the Heisenberg on the honeycomb lattice we choose a single-site unit cell,
 761 consisting of only two different tensors on the honeycomb lattice. A mapping to the square lat-
 762 tice yields a fully translationally invariant iPEPS with a local Hilbert space dimension of $p^2 = 4$.
 763 We optimize the ground states on both TN structures with $2p\chi_B^3$ and $p^2\chi_B^4$ numbers of varia-
 764 tional parameters, respectively (assuming real tensor coefficients). The model is known to be
 765 in a gapless Néel ordered phase [100–102]. Therefore, high environment bond dimensions χ_E
 766 are required to capture the large correlation lengths of the critical state. Ground state energies
 767 are reported in Fig. 29. The critical property of the ground state is already nice reflected in the
 768 significant difference between simple update MF and CTMRG expectation values. The CTMRG
 769 environments treat quantum correlations much more carefully, which leads to improved ener-
 770 gies for the infinite TN state. The VU provides lower energies than the SU with CTMRG and
 771 our results using the VU are compatible with previous results using variational iPEPS with a
 772 different CTMRG procedure [98] as well as extrapolated and thus non-variational results from
 773 the coupled cluster method [99].

774 4.3 Kagome lattice

775 The Heisenberg model on the Kagome lattice can be considered one of the most enigmatic and
 776 well studied models in the field of frustrated magnetism [104]. While a spin liquid ground




χ_B	E_0 (SU)	E_0 (VU)
2	-0.38620	-0.40454
3	-0.41786	-0.42688
4	-0.42323	-0.43038
5	-0.42866	-0.43286
6	-0.43188	-0.43451
7	-0.43313	-0.43527
8	-0.43391	-0.43552

Figure 30: Benchmarking results for the isotropic spin-1/2 Heisenberg model on the Kagome lattice. For comparison, we show the outcome obtained by extrapolated iPESS results in Ref. [103], which, to be strict, is not variational as the authors noted. Additionally, we include the result computed by exact diagonalization in Ref. [104].

777 state is well established, the actual type of ground state is still under debate with different
 778 methods supporting different states (e.g., \mathbb{Z}_2 gapped spin liquid [105, 106], $U(1)$ gapless spin
 779 liquid [103, 107]).

780 Since the ground state is known to be a spin liquid state, that does not form any magnetic
 781 ordering down to zero temperature while preserving lattice translation and rotation symmetry,
 782 we use the smallest unit cells of only three sites in our simulations. The SU then works on the
 783 three-site iPESS ansatz. The VU is performed both on the honeycomb iPESS and on a coarse-
 784 grained, fully translationally invariant iPEPS state. The number of variational parameters are
 785 hence $(3p\chi_B^2 + 2\chi_B^3)$ for the iPESS and $p^3\chi_B^4$ for the iPEPS. Again, the iPEPS state is more
 786 expressive and produces lower variational energies, that follow a smoother convergence with
 787 bond dimension χ_B , see Fig. 30. The ED energy provides a lower-bound for the energy, as
 788 argued in Sec. 4.1. Our energies are compatible with other state-of-the-art numerical methods
 789 as the extrapolated iPESS result from Ref. [103], but we would like to point out that the authors
 790 noted that their results are not variational and hence the comparison is slightly tainted. Our
 791 result showcases the purpose of variational iPEPS optimization for highly frustrated systems
 792 to obtain a real upper bound to the ground state energy.




χ_B	E_0 (SU)	E_0 (VU)
2	-0.42802	-0.43364
3	-0.43511	-0.43738
4	-0.43924	-0.43988
5	-0.43967	-0.44017
6	-0.44006	-0.44039
7	-0.44038	-0.44060

Figure 31: Benchmarking results for the isotropic spin-1/2 Heisenberg model on the square-Kagome lattice. For comparison, we include the variational Monte-Carlo results presented in Ref. [108]. Additionally, we show the extrapolated iPEPS result obtained in Ref. [109], which, to be strict, is not variational. We stress that the mean-field energies also are not variational as discussed in Sec. 4.

793 4.4 Square-Kagome lattice

794 As a third benchmark model, we simulate the Heisenberg model on the square-Kagome lattice,
 795 a lattice that has gained attention as a class of promising quantum spin liquid materials [110].
 796 It consists of corner-sharing triangles, that generate a high geometric frustration similar to the
 797 Kagome lattice. Its ground state has been found to be non-magnetic, however the existing
 798 subtle competition between different types of *valence bond crystal* (VBC) states has only been
 799 resolved recently in a TN study [109], in favor of a VBC with loop-six resonances. Simulations
 800 of the model are performed for a twelve-site checkerboard unit cell, as shown in Fig. 26.
 801 Results for the ground state energy are presented in Fig. 31. Due to the VBC ground state
 802 with a small correlation length and an energy gap in the model, the simple update MF and
 803 CTMRG energies are nearly identical. The variational update is performed on a so-called semi-
 804 PEPS structure as described in Ref. [109] and also on a coarse-grained iPEPS TN as introduced
 805 in Fig. 24, a structure that is unfeasible for SU simulations due to the large imaginary time
 806 evolution operators. Although the VU cannot significantly improve the ground state energy for
 807 the semi-PEPS ansatz, the VU on the full coarse-grained iPEPS structure improves the energies
 808 at the same bond dimension χ_B . This is connected to the larger expressivity of the coarse-
 809 grained structure.

810 Our results outperform variational Monte-Carlo simulations in Ref. [108] and are compa-



χ_B	E_0 (SU)	E_0 (VU)
2	-0.51202	-0.52675
3	-0.54181	-0.54503
4	-0.54679	-0.54867
5	-0.54756	-0.54990
6	-0.54957	-0.55147

Figure 32: Benchmarking results for the isotropic spin-1/2 Heisenberg model on the triangular lattice with an $ABC-BCA-CAB$ 3×3 unit cell structure. For comparison, we include the extrapolated, thus non-variational coupled cluster results presented in Ref. [99]. Additionally, we show the extrapolated iPESS result obtained in Ref. [111], which, to be strict, is not variational.

811 rable to state-of-the-art iPEPS results in Ref. [109]. We emphasize that the latter result is in
 812 the extrapolation, strictly speaking, not variational so that a comparison is slightly tainted.

813 4.5 Triangular lattice

814 As a last benchmark model we consider the Heisenberg model on the triangular lattice. Due to
 815 its connectivity of six, the triangular lattice exhibits a large amount of geometric frustration.
 816 The ground state is believed to be a three-sublattice 120° magnetically ordered state [112,
 817 113]. The ground state of the Heisenberg model on the triangular lattice is computed using
 818 a three-sublattice unit cell arranged in an ABC-BCA-CAB structure. The simple update data
 819 has been produced by an iPESS ansatz with the simplices sitting in the upward triangles (see
 820 Fig. 28). The VU is performed in two steps, using the converged iPESS state as input for second
 821 coarse-grained optimization run.

822 The results of our benchmark are shown in Fig. 32. In the case of the triangular lattice
 823 it generally helps to add some noise on the SU input state to reach better ground states and
 824 energies. We compare against a recent iPESS study based on the simple update [111], that
 825 predicts a zero-temperature magnetisation consistent with previous Monte Carlo studies [114]
 826 and additionally against a result obtained by the extrapolated, thus non-variational coupled
 827 cluster method [99]. We would like to point out that the iPESS result was extrapolated and
 828 is, strictly speaking, not variational.

829 4.6 Comments on excited states

830 In this work, we have primarily focused on providing a comprehensive discussion of the use of
 831 AD for the study of ground state properties of interacting quantum lattice models. It should go
 832 without saying, however, that excited states can be included in a straightforward manner. The
 833 study of excited states has first been initiated in the realm of matrix product states [115], but
 834 has later been generalized to iPEPS [116–118], allowing for constructing variational ansatzes
 835 for elementary excitations on PEPS ground states that facilitate computing gaps, dispersion
 836 relations, and spectral weights in the thermodynamic limit.

837 More recently, automatic differentiation has also found its way into the optimisation of ex-
 838 cited states [42]. The central idea is to construct the excited state with momentum $\vec{k} = (k_x, k_y)$
 839 as a superposition of the ground state vector, perturbed by a single tensor B at position $\vec{x} = (x, y)$
 840 and appropriate phase factors according to

$$|\phi(B)_{\vec{k}}\rangle = \sum_{\vec{x}} e^{i\vec{k}\vec{x}} |\phi(B)_{\vec{x}}\rangle. \quad (24)$$

841 The coefficients of tensor B are then determined by energy minimisation of the excited state,
 842 for which AD can again be used [42, 119]. In contrast to the regular ground state optimisation,
 843 here the CTMRG routine must be extended to include the appropriate phase factors in the
 844 directional absorption. Moreover, instead of only eight environment tensors per iPEPS tensor
 845 in the unit cell, the action of B , B^\dagger and the product of B and B^\dagger has to be tracked in three
 846 additional sets of eight tensors.

847 The excited state approach can be directly extended to different lattice geometries. To
 848 this end, we have to generalize the absorption of iPEPS tensors (growing the CTMRG transfer
 849 tensors T_1 , T_2 , T_3 and T_4) to include the basis of the lattice, respecting relative phase factors
 850 of the basis vectors. Depending on the actual structure of the basis, a separate tensor B_n is
 851 chosen as a perturbation for each of the basis site. Our implementation already contains the
 852 main building blocks of a robust and flexible CTMRG routine, calculation of gradients using AD
 853 at the fixed-point and minimisation of an energy cost function. The extension of the framework
 854 to include excited states is therefore natural. It is planned as a future feature.

855 4.7 Comments on fermionic systems

856 As a final comment we stress that for clarity and to be concise, we have focused in our pre-
 857 sentation on quantum spin models. It should be clear, however, that the machinery developed
 858 here readily carries over to the study of *interacting fermionic systems*, with little modifications.
 859 Naively, one might think that the simulation of two-dimensional fermionic models is marred
 860 by substantial overheads that emerge when invoking a spin-to-fermion mapping. This is, how-
 861 ever, not the case, and the respective book-keeping of the signs can be done with negligible
 862 overhead [120, 121]. On the formal level, such tensor networks involve a particular choice
 863 of what is called a spin structure [122, 123]. Practically speaking, one can modify much of
 864 the bosonic code for PEPS to the fermionic setting, readily incorporating the relevant signs
 865 to capture interacting fermions, in what is called *fermionic PEPS* [120, 124, 125]. This insight
 866 is important as some of the most compelling test cases of interacting quantum many-body
 867 systems are of a fermionic nature.

868 5 Conclusion and prospects

869 In this review we present a comprehensive introduction into automatic differentiation in the
 870 context of two-dimensional tensor networks, leading to the recently emerging variational

871 iPEPS framework for ground state optimization. We provide implementation details and dis-
 872 cuss obstacles that arise in practice, as well as techniques to mitigate these. At the same time,
 873 we coherently present ideas that have to date only been mentioned in a fragmented fashion
 874 in the literature. We hope that the present work can serve as a useful reference and review in
 875 the variational study of 2d tensor networks.

876 This work accompanies the variational iPEPS library *variPEPS*, a comprehensive and ver-
 877 satile code base for optimizing iPEPS in a general setting. We expect this library to be a helpful
 878 tool for performing state-of-the-art tensor network analyses for a wide range of physical mod-
 879 els, featuring multiple two-dimensional lattices. The library is designed to be extended with
 880 additional simulation techniques based on automatic differentiation, such as excited states and
 881 structure factors.

882 The *variPEPS* library is publicly available in both a Julia and a Python version on GitHub [56],
 883 with stable references in the corresponding Zenodo repositories [57, 58].

884 5.1 CO₂-emissions table

885 For the sake of completeness and for promoting carbon footprint awareness, we display an
 886 estimated lower bound of the carbon emissions generated during the course of this work in
 887 Table 2.

Numerical simulations	
Total Kernel Hours [h]	≥ 255276
Thermal Design Power Per Kernel [W]	12
Total Energy Consumption Simulations [kWh]	≥ 3063
Average Emission Of CO ₂ In Germany [kg/kWh]	0.441
Total CO ₂ -Emission For Numerical Simulations [kg]	≥ 1351
Were The Emissions Offset?	Yes
Air Travel	
Total CO ₂ -Emission For Air Travel [kg]	924
Were The Emissions Offset?	Yes
Total CO ₂ -Emission [kg]	≥ 2275

Table 2: Summary of the estimated lower bound of the carbon cost generated during the development of this work. The estimations have been calculated using the examples of the Scientific CO₂nduct project [126] and include the costs of the numerical calculations and air travel for collaborations.

888 Acknowledgements

889 We acknowledge inspiring discussions with Ji-Yao Chen, Andreas Haller, Juraj Hasik, Augus-
 890 tine Kshetrimayum, Alexander Nietner and Niklas Tausendpfund. We would like to particu-
 891 larly thank Boris Ponsioen, who has for shared valuable insights, and Frederik Wilde, who has
 892 helped us to get the details of the automatic differentiation and the custom fixed-point deriva-
 893 tive correct. We would like to thank the ZEDV (IT support) of the physics department, Freie
 894 Universität Berlin, for computing time and their technical support, particularly we thank Jörg
 895 Behrmann and Jens Dreger.

896 For the Python version we thank the developers of the JAX framework [127, 128] for their
 897 work to provide an AD framework with optimized numerical operations and their technical
 898 support during the development of this work. We also acknowledge the use of the TensorKit

899 package [129] in the Julia version of the code and wish to advertise the open source libraries
900 of the Quantum Ghent group in this context [130]. We make use of the Zygote [131] package
901 for AD in the Julia programming language.

902 The work has been discussed and refined during the workshop “Tensor Networks: Math-
903 ematical Structures and Novel Algorithms (2022)” at the Erwin Schrödinger International In-
904 stitute for Mathematics and Physics in Vienna and the workshop “Entanglement in Strongly
905 Correlated Systems (2023)” at the Centro de Ciencias de Benasque Pedro Pascual. We thank
906 the organizers for their hospitality and work.

907 **Funding information** E. L. W. thanks the Studienstiftung des deutschen Volkes for support.
908 This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research
909 Foundation) under the project number 277101999 – CRC 183 (project B01), for which this con-
910 stitutes an inter-node publication involving both Cologne and Berlin, and the BMBF (MUNIQ-
911 Atoms, FermiQP). It has also received funding from the Cluster of Excellence MATH+ and from
912 the Quantum Flagship (PasQuans2).

913 The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-
914 centre.eu) for funding this project by providing computing time through the John von Neu-
915 mann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at the Jülich Su-
916 percomputing Centre (JSC) (Grant NeTeNeSyQuMa) and the FZ Jülich for JURECA (institute
917 project PGI-8) [132].

918 Appendix: Background on automatic differentiation

919 A Adjoint functions and variables

920 In the literature it is common to use so called *adjoint functions* and *adjoint variables* when
 921 using backwards-mode AD. These adjoint functions map the adjoint variables onto each other,
 922 as in Eq. (15) when building up the gradient. In this section, we will briefly introduce the
 923 basic notation of adjoint functions and variables following Ref. [133]. Explicit constructions
 924 of adjoint functions, which are vector-Jacobian-products in the practical implementation, for
 925 a large number of useful operations including those for the iPEPS use-case can be found in
 926 Refs. [133–135].

927 As an example throughout this section, we consider the function h , composed out of two
 928 primitive functions h_1 and h_2 which are concatenated as

$$\begin{aligned} h &= h_2 \circ h_1, \\ h_1 &: M_{n \times n} \times M_{n \times n} \rightarrow M_{n \times n}, \\ h_2 &: M_{n \times n} \rightarrow \mathbb{R}, \end{aligned} \quad (\text{A.1})$$

929 with variables $(A, B) \in M_{n \times n} \times M_{n \times n}$, $C \in M_{n \times n}$ and $x \in \mathbb{R}$. We start by examining the differ-
 930 ential of the output variable x

$$dx = \frac{\partial h_2}{\partial C} dC =: \sum_{i,j} \bar{C}_{i,j} dC_{i,j} = \text{Tr}(\bar{C}^\top dC). \quad (\text{A.2})$$

931 In the first equation, we have suppressed the sum over the indices of C . Eq. (A.2) defines the
 932 adjoint variable \bar{C} of C . We see that the adjoint variable \bar{C} is the derivative of the scalar output
 933 of the function h_2 w.r.t. C . Thus, for the case of a scalar output the variable C and the adjoint
 934 variable \bar{C} have the same dimension. Now, in order to get the gradient ∇h we are interested
 935 in the derivative of the output w.r.t. the input variables (A, B) . To this end we consider the
 936 differential of the intermediate variable

$$dC = \frac{\partial h_1}{\partial A} dA + \frac{\partial h_1}{\partial B} dB. \quad (\text{A.3})$$

937 Inserting this into Eq. (A.2), we obtain

$$dx = \text{Tr} \left(\underbrace{\bar{C}^\top \frac{\partial h_1}{\partial A}}_{\bar{A}^\top} dA \right) + \text{Tr} \left(\underbrace{\bar{C}^\top \frac{\partial h_1}{\partial B}}_{\bar{B}^\top} dB \right). \quad (\text{A.4})$$

938 Here we have already implicitly used the adjoint function \bar{h}_1 that maps the adjoint variable \bar{C}
 939 to the adjoint variables \bar{A} and \bar{B} according to

$$\bar{h}_1 : \bar{C}^\top \mapsto (\bar{A}^\top, \bar{B}^\top)^\top = \left(\bar{C}^\top \frac{\partial h_1}{\partial A}, \bar{C}^\top \frac{\partial h_1}{\partial B} \right)^\top. \quad (\text{A.5})$$

940 Given the fact that we are dealing with a scalar output variable x , we recall that \bar{C} can be
 941 considered a vector, such that the adjoint function is a vector-Jacobian-product (vJP). We can
 942 see that the this mapping of the adjoint variables with adjoint functions eventually produces
 943 the gradient

$$\begin{aligned} \nabla h = (\bar{A}, \bar{B}) &= \left(\frac{\partial h_1}{\partial A} \bar{C}, \frac{\partial h_1}{\partial B} \bar{C} \right) \\ &= \left(\frac{\partial h_1}{\partial A} \frac{\partial h_2}{\partial C}, \frac{\partial h_1}{\partial B} \frac{\partial h_2}{\partial C} \right). \end{aligned} \quad (\text{A.6})$$

944 B Automatic differentiation for complex variables

945 Some extra attention has to be given to the case in which the primitive functions are complex
 946 valued. This is because not all functions one might want to consider are complex-differentiable
 947 (holomorphic) and as such the derivative depends on the direction we move in the complex
 948 plane when taking the limit for the derivative. In such a case one needs to resort to the calculus
 949 of two sets of independent real variables. For a generic function $f : \mathbb{C} \rightarrow \mathbb{C}$ this can be done
 950 by treating x and y in $z = x + iy$ as independent variables or alternatively, by choosing z
 951 and z^* and making use of Wirtinger calculus. However we should also note that in the iPEPS
 952 use case we deal with a function $E : \mathbb{C}^n \rightarrow \mathbb{R}$, which removes the necessity to think about
 953 holomorphism.

954 C The implicit function theorem and its use at the CTMRG fixed- 955 point

956 In this section, we are going to present an alternative approach to taking the derivative of the
 957 energy function by utilizing the fixed point of the CTMRG procedure. To this end, we can
 958 make use of the implicit function theorem [136] to calculate the derivative of the full fixed-
 959 point routine. Our discussion will follow the description of Refs. [137, 138]. Differentiating
 960 Eq. (16) on both sides we end up with

$$\partial_A e^*(A) = \partial_A c(A, e^*) + \partial_{e^*} c(A, e^*) \partial_A e^*(A). \quad (\text{C.1})$$

961 Introducing the shorthand writing for the Jacobians $L = \partial_A c(A, e^*(A))$ and $K = \partial_{e^*} c(A, e^*(A))$
 962 and rearranging the equation we find

$$\begin{aligned} \partial_A e^*(A) &= (L + K \partial_A e^*(A)) \\ &= \left(\sum_{n=0}^{\infty} K^n \right) L = (\mathbb{1} - K)^{-1} L. \end{aligned} \quad (\text{C.2})$$

963 As discussed in Appendix A, we aim at finding the adjoint function of the CTMRG iteration at
 964 the fixed point, which is a *vector-Jacobian product* (vJP) $\mathbf{v}^\top \partial_A e^*(A)$. Inserting Eq. (C.2) yields

$$\mathbf{v}^\top \partial e^*(A) = \mathbf{v}^\top (\mathbb{1} - K)^{-1} L = \mathbf{w}^\top L, \quad (\text{C.3})$$

965 where we have introduced $\mathbf{w}^\top := \mathbf{v}^\top (\mathbb{1} - K)^{-1}$. The second equality in the equation above can
 966 be rearranged into another fixed-point equation

$$\mathbf{w}^\top = \mathbf{v}^\top + \mathbf{w}^\top K. \quad (\text{C.4})$$

967 Here $\mathbf{w}^\top K$ is another vJP but this time only dependent of the derivative of a single absorption
 968 step evaluated at the fixed-point of the CTMRG routine. Solving Eq. (C.4) we can find \mathbf{w}^\top
 969 to calculate the vJP of the CTMRG routine from Eq. (C.3). In the end we reduced the naive
 970 effort of unrolling the fixed-point iterations to just calculate the derivative of a single CTMRG
 971 iteration and another fixed-point iteration which both are much less memory intensive.

972 D Automatic differentiation in the language of differential geom- 973 etry

974 In order to unify the different frameworks for thinking about forward- and backwards-mode
 975 AD, we will briefly introduce a mathematical notation for AD. It also serves to give some

976 more precise meaning to the terms “push-forward” and “pullback”, that are sometimes used in
 977 forward- and backwards-mode AD discussions, respectively. For this we first recall the general
 978 concept of a push-forward and a pullback for the simple case of functions and distributions.
 979 Imagine two functions $f : M \rightarrow N$ and $g : N \rightarrow \mathbb{R}$. The *pullback* of g along f allows us to
 980 construct a function $f^*g : M \rightarrow \mathbb{R}$ for which the domain of the function g is “pulled back” to
 981 the domain of the function f . This is done by a simple concatenation of f and g

$$f^*g(\underbrace{m}_{\in M}) = (g \circ f)(m) = g(f(m)). \tag{D.1}$$

982 This construction can now be used to define a *push-forward* on the dual objects of the functions
 983 under integration. These dual objects are distributions. With a distribution, we can integrate
 984 a function

$$\int_M \bullet \mu : \mathcal{F}(M) \rightarrow \mathbb{R},$$

$$f \mapsto \int_M f \mu, \tag{D.2}$$

985 where $\mathcal{F}(M)$ are just the functions on M and μ is the distribution. Given such a distribution
 986 on M we can now integrate functions on M . The push-forward $f_*\mu$ of μ allows us to integrate
 987 functions on N by defining a distribution that is “pushed forward” to N . This works as

$$\int_N h(f_*\mu) =: \int_M (f^*h)\mu, \tag{D.3}$$

988 where h is a function on N .

989 This type of construction for the pullback and push-forward generalizes to many mathe-
 990 matical objects that have a *pairing dual*. The relevant mathematical objects for AD are the
 991 derivative $\partial/\partial x_i$ and its pairing dual, the differential dx_i .

992 It might be useful, beyond the conceptual clarity of this notation, to look at AD in this
 993 way because one can easily imagine situations where the intermediate data of a function is
 994 restricted by constraints such that the “data-space” becomes geometrically non-trivial. An
 995 example could be vectors in \mathbb{R}^n restricted to unit length or matrices in $M_{n,m}$ restricted to be
 996 unitary. We note that an optimisation in these situations requires some additional concepts,
 997 like finding a path on the given space from a tangent vector. This requires some extra care and
 998 is not discussed here.

We now introduce the mathematical notation that we need in order to talk about AD in
 this language. We will not be particularly rigorous in this endeavour and leave out all details
 that are not explicitly needed. We start with a manifold M on which we can consider points
 $p \in M$, as well as functions $f : M \rightarrow \mathbb{R}$. For each point $p \in M$ we can define a vector space
 T_pM (call it the tangent-space at p) of tangent-vectors at that point. The elements in T_pM act
 like derivatives on functions on M

$$\text{e.g.: } \frac{\partial}{\partial x_i} = e_i \in T_pM, \quad \frac{\partial}{\partial x_i}(f) = \frac{\partial f}{\partial x_i}.$$

Here we have assumed that we have equipped the manifold M with coordinates via a chart
 $\phi : M \rightarrow \mathbb{R}^m$ around the point p , where $m = \dim(M)$. Our tangent-space T_pM has dimension
 m and we can choose a canonical basis

$$\left\{ \frac{\partial}{\partial x_1}, \dots, \frac{\partial}{\partial x_m} \right\} = \{e_1, \dots, e_m\}.$$

One further defines the dual vector space T_p^*M of the tangent vector space, called cotangent-space. This cotangent-space contains the dual vectors to the derivatives $\frac{\partial}{\partial x_i}$. These cotangent vectors from the cotangent-space are the differentials dx_i . The cotangent-space also has dimension m and we can choose the canonical basis

$$\{dx_1, \dots, dx_m\}.$$

999 Obviously, given the canonical basis for the tangent-space and cotangent-space we can expand
1000 arbitrary vectors in these spaces in the basis. Take $v \in T_pM$ and $df \in T_p^*M$ we can expand as

$$v = \sum_i v_i \frac{\partial}{\partial x_i} = \sum_i v_i e_i, \quad (\text{D.4})$$

$$df = \sum_i \frac{\partial f}{\partial x_i} dx_i. \quad (\text{D.5})$$

1001 We have a pairing between the derivatives that live in the tangent-space T_pM and the differ-
1002 entials that live in T_p^*M as

$$dx_j \left(\frac{\partial}{\partial x_i} \right) := \frac{\partial x_j}{\partial x_i} = \delta_{i,j}. \quad (\text{D.6})$$

1003 Note that by this pairing relation we see that tangent and cotangent vectors are “pairing duals”
1004 and we can use an analogous construction for pullbacks and push-forwards as we did for
1005 functions and distribution above. Since T_pM and T_p^*M are isomorphic, we can introduce a
1006 correspondence transformation between the canonical basis of the two spaces

$$\bullet^b : T_pM \rightarrow T_p^*M, \quad e_i \mapsto dx_i = e_i^b, \quad (\text{D.7})$$

$$\bullet^\sharp : T_p^*M \rightarrow T_pM, \quad dx_i \mapsto e_i = dx_i^\sharp. \quad (\text{D.8})$$

1007 We now have assembled all necessary tools to formulate what a “gradient” is in this language.
1008 It is given by

$$\nabla f := (df)^\sharp, \quad (\text{D.9})$$

1009 which matches the common formula

$$\begin{aligned} \nabla f &= \left(\sum_i \frac{\partial f}{\partial x_i} dx_i \right)^\sharp = \sum_i \frac{\partial f}{\partial x_i} e_i \\ &= \left(\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_m} \right), \end{aligned} \quad (\text{D.10})$$

1010 where we have taken e_i just as the i -th unit vector of T_pM .

1011 Now it is easy to construct the pullbacks and push-forwards in this context analogous to our
1012 treatment of functions and distributions. For this we start from manifolds M and N with points
1013 $p \in M$ and $q \in N$, and with the two functions $f : M \rightarrow N$ and $g : N \rightarrow \mathbb{R}$. We can consider a
1014 differential $dg \in T_q^*N$ which we want to “pull back” along the function f and associate it with
1015 and element of $T_{f^{-1}(q)}^*M$, where $f^{-1}(q) \in M$. We do this with the familiar definition

$$\underbrace{f^* dg}_{\in T_{f^{-1}(q)}^*M} := d(g \circ f) \quad (\text{D.11})$$

1016 which uses a concatenation of f and g just as in the first example. For a tangible example
1017 consider $g = x_i$ to be a coordinate function. We then get $f^* dx_i = d(x_i \circ f) = d(f_i)$. As

1018 before the push-forward can be defined via the pullback just as we had done for functions and
 1019 distributions. In this case, we start with a tangent vector $\frac{\partial}{\partial x_i}$ in $T_p M$ and want to “push it
 1020 forward” along f into $T_{f(p)} N$. This works as

$$\underbrace{\left(f_* \left(\frac{\partial}{\partial x_i} \right) \right)}_{\in T_{f(p)} N} (g) := \frac{\partial}{\partial x_i} (f^* g) = \frac{\partial}{\partial x_i} (g \circ f). \quad (\text{D.12})$$

1021 Now that we are equipped with the pullback and push-forward of differentials and derivatives
 1022 we see how the gradient is calculated in the forward- and backward-mode AD. For this we will
 1023 go back to our neat example from Sec. 2.4 and slightly generalize. Say, we would like to take
 1024 the gradient ∇E of a function that is composed of three primitive functions $E = f_3 \circ f_2 \circ f_1$.
 1025 We say these primitive functions map between manifolds

$$E : M_1 \xrightarrow{f_1} M_2 \xrightarrow{f_2} M_3 \xrightarrow{f_3} \mathbb{R}. \quad (\text{D.13})$$

1026 Lets first look at what happens when we build the gradient using backwards-mode AD. In this
 1027 case we start with the differential df_3 of the last primary function of E . This differential lives
 1028 in $T_k^* M_3$, where $k \in M_3$ is a point in M_3 . We can now use the pullback along the functions f_2
 1029 and then f_1 to pull back this differential to M_1

$$f_1^*(f_2^*(df_3)) \xleftarrow{\text{pullback}} f_2^*(df_3) \xleftarrow{\text{pullback}} df_3. \quad (\text{D.14})$$

1030 With the definitions above we see that in this way we construct the gradient

$$f_1^*(f_2^*(df_3)) = f_1^*((d(f_3 \circ f_2))) = d(f_3 \circ f_2 \circ f_1) = dE. \quad (\text{D.15})$$

1031 With our identification between tangent and cotangent vectors we finalize to $\nabla E = (dE)^\sharp$. If
 1032 we express the differential that we start from df_3 in coordinates, we straightforwardly obtain
 1033 the product of Jacobians as a result for the gradient. This also establishes the connection to
 1034 the adjoint functions we talked about in the previous section and the vector-Jacobian product
 1035 as discussed in Sec. 2.4.

1036 In the case of forward-mode AD we start from a tangent vector $\frac{\partial}{\partial x_i}$, which lives in $T_l M_1$,
 1037 where $l \in M_1$ is a point in M_1 . We can now push this tangent vector forward into a tangent
 1038 space of M_3 with successive push-forwards along f_1 followed by f_2

$$\frac{\partial}{\partial x_i} \xrightarrow{\text{push-forward}} f_{1*} \left(\frac{\partial}{\partial x_i} \right) \xrightarrow{\text{push-forward}} f_{2*} \left(f_{1*} \left(\frac{\partial}{\partial x_i} \right) \right). \quad (\text{D.16})$$

1039 With the definitions for the push-forward we see that the gradient we obtain in this way is
 1040 given by

$$\begin{aligned} \sum_i f_{2*} \left(f_{1*} \left(\frac{\partial}{\partial x_i} \right) \right) (f_3) e_i &= \sum_i f_{1*} \left(\frac{\partial}{\partial x_i} \right) (f_3 \circ f_2) e_i \\ &= \sum_i \frac{\partial}{\partial x_i} \underbrace{(f_3 \circ f_2 \circ f_1)}_{=E} e_i \\ &= \nabla E. \end{aligned} \quad (\text{D.17})$$

References

- 1041
- 1042 [1] T. Nishino, *DMRG homepage* (2020), <http://quattro.phys.sci.kobe-u.ac.jp/dmrg/>
1043 [condmat91.html](http://quattro.phys.sci.kobe-u.ac.jp/dmrg/condmat91.html).
- 1044 [2] H. A. Kramers and G. H. Wannier, *Statistics of the two-dimensional ferromagnet. part ii*,
1045 Phys. Rev. **60**, 263 (1941), doi:[10.1103/PhysRev.60.263](https://doi.org/10.1103/PhysRev.60.263).
- 1046 [3] R. Baxter, *Dimers on a rectangular lattice*, J. Math. Phys. **9**, 650–654 (1968),
1047 doi:[10.1063/1.1664623](https://doi.org/10.1063/1.1664623).
- 1048 [4] S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev.
1049 Lett. **69**, 2863 (1992), doi:[10.1103/PhysRevLett.69.2863](https://doi.org/10.1103/PhysRevLett.69.2863).
- 1050 [5] U. Schollwöck, *The density-matrix renormalization group in the age of matrix product*
1051 *states*, Ann. Phys. **326**, 96 (2011), doi:[10.1016/j.aop.2010.09.012](https://doi.org/10.1016/j.aop.2010.09.012).
- 1052 [6] S. Östlund and S. Rommer, *Thermodynamic limit of density matrix renormalization*,
1053 Phys. Rev. Lett. **75**, 3537 (1995), doi:[10.1103/PhysRevLett.75.3537](https://doi.org/10.1103/PhysRevLett.75.3537).
- 1054 [7] J. Dukelsky, M. A. Martin-Delgado, T. Nishino and G. Sierra, *Equivalence of the vari-*
1055 *ational matrix product method and the density matrix renormalization group applied to*
1056 *spin chains*, Europhys. Lett. **43**, 457 (1998), doi:[10.1209/epl/i1998-00381-x](https://doi.org/10.1209/epl/i1998-00381-x).
- 1057 [8] D. Perez-Garcia, F. Verstraete, M. Wolf and J. Cirac, *Matrix product state representations*,
1058 Quant. Inf. Comp. **7**, 401 (2007), doi:[10.26421/QIC7.5-6-1](https://doi.org/10.26421/QIC7.5-6-1).
- 1059 [9] M. Fannes, B. Nachtergaele and R. Werner, *Finitely correlated states on quantum spin*
1060 *chains*, Commun. Math. Phys. **144**, 443 (1992), doi:[10.1007/BF02099178](https://doi.org/10.1007/BF02099178).
- 1061 [10] C. Lubich, I. V. Oseledets and B. Vandereycken, *Time integration of tensor trains*, SIAM
1062 J. Num. An. **53**(2), 917 (2015), doi:[10.1137/140976546](https://doi.org/10.1137/140976546).
- 1063 [11] R. Orús, *A practical introduction to tensor networks: Matrix product states and projected*
1064 *entangled pair states*, Ann. Phys. **349**, 117 (2014), doi:[10.1016/j.aop.2014.06.013](https://doi.org/10.1016/j.aop.2014.06.013).
- 1065 [12] J. C. Bridgeman and C. T. Chubb, *Hand-waving and interpretive dance: An introduc-*
1066 *tory course on tensor networks*, J. Phys. A **50**, 223001 (2017), doi:[10.1088/1751-](https://doi.org/10.1088/1751-8121/aa6dc3)
1067 [8121/aa6dc3](https://doi.org/10.1088/1751-8121/aa6dc3).
- 1068 [13] J. I. Cirac, D. Pérez-García, N. Schuch and F. Verstraete, *Matrix product states and*
1069 *projected entangled pair states: Concepts, symmetries, theorems*, Rev. Mod. Phys. **93**,
1070 045003 (2021), doi:[10.1103/RevModPhys.93.045003](https://doi.org/10.1103/RevModPhys.93.045003).
- 1071 [14] F. Verstraete, J. I. Cirac and V. Murg, *Matrix product states, projected entangled pair*
1072 *states, and variational renormalization group methods for quantum spin systems*, Adv.
1073 Phys. **57**, 143 (2008), doi:[10.1080/14789940801912366](https://doi.org/10.1080/14789940801912366).
- 1074 [15] J. Eisert, M. Cramer and M. B. Plenio, *Area laws for the entanglement entropy*, Rev. Mod.
1075 Phys. **82**, 277 (2010), doi:[10.1103/RevModPhys.82.277](https://doi.org/10.1103/RevModPhys.82.277).
- 1076 [16] F. Verstraete and J. I. Cirac, *Renormalization algorithms for Quantum-Many Body Systems*
1077 *in two and higher dimensions* (2004), <https://arxiv.org/abs/cond-mat/0407066>.
- 1078 [17] D. Perez-Garcia, F. Verstraete, M. M. Wolf and J. I. Cirac, *PEPS as unique ground states*
1079 *of local Hamiltonians*, Quant. Inf. Comp. **8**, 650 (2008), doi:[10.26421/QIC8.6-7-6](https://doi.org/10.26421/QIC8.6-7-6).

- 1080 [18] N. Schuch, D. Perez-Garcia and I. Cirac, *Classifying quantum phases using matrix*
1081 *product states and projected entangled pair states*, Phys. Rev. B **84**, 165139 (2011),
1082 doi:[10.1103/PhysRevB.84.165139](https://doi.org/10.1103/PhysRevB.84.165139).
- 1083 [19] N. Schuch, M. M. Wolf, F. Verstraete and J. I. Cirac, *Computational com-*
1084 *plexity of projected entangled pair states*, Phys. Rev. Lett. **98**, 140506 (2007),
1085 doi:[10.1103/PhysRevLett.98.140506](https://doi.org/10.1103/PhysRevLett.98.140506).
- 1086 [20] J. Haferkamp, D. Hangleiter, J. Eisert and M. Gluza, *Contracting projected en-*
1087 *tangled pair states is average-case hard*, Phys. Rev. Research **2**, 013010 (2020),
1088 doi:[10.1103/PhysRevResearch.2.013010](https://doi.org/10.1103/PhysRevResearch.2.013010).
- 1089 [21] M. Schwarz, O. Buerschaper and J. Eisert, *Approximating local observ-*
1090 *ables on projected entangled pair states*, Phys. Rev. A **95**, 060102 (2017),
1091 doi:[10.1103/PhysRevA.95.060102](https://doi.org/10.1103/PhysRevA.95.060102).
- 1092 [22] J. Jordan, R. Orús, G. Vidal, F. Verstraete and J. I. Cirac, *Classical simulation of infinite-*
1093 *size quantum lattice systems in two spatial dimensions*, Phys. Rev. Lett. **101**, 250602
1094 (2008), doi:[10.1103/PhysRevLett.101.250602](https://doi.org/10.1103/PhysRevLett.101.250602).
- 1095 [23] H. N. Phien, J. A. Bengua, H. D. Tuan, P. Corboz and R. Orús, *Infinite projected entangled*
1096 *pair states algorithm improved: Fast full update and gauge fixing*, Phys. Rev. B **92**, 035142
1097 (2015), doi:[10.1103/PhysRevB.92.035142](https://doi.org/10.1103/PhysRevB.92.035142).
- 1098 [24] R. Orús and G. Vidal, *Simulation of two-dimensional quantum systems on an infinite*
1099 *lattice revisited: Corner transfer matrix for tensor contraction*, Phys. Rev. B **80**, 094403
1100 (2009), doi:[10.1103/PhysRevB.80.094403](https://doi.org/10.1103/PhysRevB.80.094403).
- 1101 [25] T. Nishino and K. Okunishi, *Corner transfer matrix renormalization group method*, J.
1102 Phys. Soci. Jap, **65**(4), 891 (1996), doi:[10.1143/JPSJ.65.891](https://doi.org/10.1143/JPSJ.65.891).
- 1103 [26] T. Nishino and K. Okunishi, *Corner transfer matrix algorithm for classical renormalization*
1104 *group*, J. Phys. Soc. Jap. **66**(10), 3040 (1997), doi:[10.1143/JPSJ.66.3040](https://doi.org/10.1143/JPSJ.66.3040).
- 1105 [27] M. Levin and C. P. Nave, *Tensor renormalization group approach to two-*
1106 *dimensional classical lattice models*, Phys. Rev. Lett. **99**, 120601 (2007),
1107 doi:[10.1103/PhysRevLett.99.120601](https://doi.org/10.1103/PhysRevLett.99.120601).
- 1108 [28] Z. Y. Xie, H. C. Jiang, Q. N. Chen, Z. Y. Weng and T. Xiang, *Second*
1109 *renormalization of tensor-network states*, Phys. Rev. Lett. **103**, 160601 (2009),
1110 doi:[10.1103/PhysRevLett.103.160601](https://doi.org/10.1103/PhysRevLett.103.160601).
- 1111 [29] H. H. Zhao, Z. Y. Xie, Q. N. Chen, Z. C. Wei, J. W. Cai and T. Xiang,
1112 *Renormalization of tensor-network states*, Phys. Rev. B **81**, 174411 (2010),
1113 doi:[10.1103/PhysRevB.81.174411](https://doi.org/10.1103/PhysRevB.81.174411).
- 1114 [30] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang and T. Xiang, *Coarse-graining renormal-*
1115 *ization by higher-order singular value decomposition*, Phys. Rev. B **86**, 045139 (2012),
1116 doi:[10.1103/PhysRevB.86.045139](https://doi.org/10.1103/PhysRevB.86.045139).
- 1117 [31] V. Zauner-Stauber, L. Vanderstraeten, M. T. Fishman, F. Verstraete and J. Haegeman,
1118 *Variational optimization algorithms for uniform matrix product states*, Phys. Rev. B **97**,
1119 045145 (2018), doi:[10.1103/PhysRevB.97.045145](https://doi.org/10.1103/PhysRevB.97.045145).

- 1120 [32] A. Nietner, B. Vanhecke, F. Verstraete, J. Eisert and L. Vanderstraeten, *Efficient varia-*
1121 *tional contraction of two-dimensional tensor networks with a non-trivial unit cell*, *Quan-*
1122 *tum* **4**, 328 (2020), doi:[10.48550/arXiv.2003.01142](https://doi.org/10.48550/arXiv.2003.01142).
- 1123 [33] C. M. Dawson, J. Eisert and T. J. Osborne, *Unifying variational methods for*
1124 *simulating quantum many-body systems*, *Phys. Rev. Lett.* **100**, 130501 (2008),
1125 doi:[10.1103/PhysRevLett.100.130501](https://doi.org/10.1103/PhysRevLett.100.130501).
- 1126 [34] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde and F. Verstraete,
1127 *Time-dependent variational principle for quantum lattices*, *Phys. Rev. Lett.* **107**, 070601
1128 (2011), doi:[10.1103/PhysRevLett.107.070601](https://doi.org/10.1103/PhysRevLett.107.070601).
- 1129 [35] P. Corboz, *Variational optimization with infinite projected entangled-pair states*, *Phys.*
1130 *Rev. B* **94**, 035133 (2016), doi:[10.1103/PhysRevB.94.035133](https://doi.org/10.1103/PhysRevB.94.035133).
- 1131 [36] L. Vanderstraeten, J. Haegeman, P. Corboz and F. Verstraete, *Gradient methods for vari-*
1132 *ational optimization of projected entangled-pair states*, *Phys. Rev. B* **94**, 155123 (2016),
1133 doi:[10.1103/PhysRevB.94.155123](https://doi.org/10.1103/PhysRevB.94.155123).
- 1134 [37] H.-J. Liao, J.-G. Liu, L. Wang and T. Xiang, *Differentiable programming tensor networks*,
1135 *Phys. Rev. X* **9**, 031041 (2019), doi:[10.1103/PhysRevX.9.031041](https://doi.org/10.1103/PhysRevX.9.031041).
- 1136 [38] C. Hubig, *Use and implementation of autodifferentiation in tensor network methods with*
1137 *complex scalars* (2019), <https://arxiv.org/abs/1907.13422>.
- 1138 [39] W.-L. Tu, H.-K. Wu, N. Schuch, N. Kawashima and J.-Y. Chen, *Generating func-*
1139 *tion for tensor network diagrammatic summation*, *Phys. Rev. B* **103**, 205155 (2021),
1140 doi:[10.1103/PhysRevB.103.205155](https://doi.org/10.1103/PhysRevB.103.205155).
- 1141 [40] J. Hasik, D. Poilblanc and F. Becca, *Investigation of the Néel phase of the frustrated*
1142 *Heisenberg antiferromagnet by differentiable symmetric tensor networks*, *SciPost Phys.*
1143 **10**, 012 (2021), doi:[10.21468/SciPostPhys.10.1.012](https://doi.org/10.21468/SciPostPhys.10.1.012).
- 1144 [41] J. Hasik, M. Van Damme, D. Poilblanc and L. Vanderstraeten, *Simulating chiral spin*
1145 *liquids with projected entangled-pair states*, *Phys. Rev. Lett.* **129**, 177201 (2022),
1146 doi:[10.1103/PhysRevLett.129.177201](https://doi.org/10.1103/PhysRevLett.129.177201).
- 1147 [42] B. Ponsioen, F. F. Assaad and P. Corboz, *Automatic differentiation applied to*
1148 *excitations with Projected Entangled Pair States*, *SciPost Phys.* **12**, 6 (2022),
1149 doi:[10.21468/SciPostPhys.12.1.006](https://doi.org/10.21468/SciPostPhys.12.1.006).
- 1150 [43] M. J. O'Rourke and G. K.-L. Chan, *Entanglement in the quantum phases of an unfrustrated*
1151 *rydberg atom array* (2022), <https://arxiv.org/abs/2201.03189>.
- 1152 [44] I. V. Lukin and A. G. Sotnikov, *Variational optimization of tensor-network states with*
1153 *the honeycomb-lattice corner transfer matrix*, *Phys. Rev. B* **107**, 054424 (2023),
1154 doi:[10.1103/PhysRevB.107.054424](https://doi.org/10.1103/PhysRevB.107.054424).
- 1155 [45] F. Wilde, A. Kshetrimayum, I. Roth, D. Hangleiter, R. Sweke and J. Eisert, *Scalably*
1156 *learning quantum many-body Hamiltonians from dynamical data* (2022), [https://arxiv.](https://arxiv.org/abs/2209.14328)
1157 [org/abs/2209.14328](https://arxiv.org/abs/2209.14328).
- 1158 [46] I. V. Lukin and A. G. Sotnikov, *Variational optimization of tensor-network states with*
1159 *the honeycomb-lattice corner transfer matrix*, *Phys. Rev. B* **107**, 054424 (2023),
1160 doi:[10.1103/PhysRevB.107.054424](https://doi.org/10.1103/PhysRevB.107.054424).

- 1161 [47] X.-Y. Zhang, S. Liang, H.-J. Liao, W. Li and L. Wang, *Differentiable program-*
1162 *ming tensor networks for kitaev magnets*, Phys. Rev. B **108**, 085103 (2023),
1163 doi:[10.1103/PhysRevB.108.085103](https://doi.org/10.1103/PhysRevB.108.085103).
- 1164 [48] F. Ferrari, S. Niu, J. Hasik, Y. Iqbal, D. Poilblanc and F. Becca, *Static and dynamical*
1165 *signatures of Dzyaloshinskii-Moriya interactions in the Heisenberg model on the kagome*
1166 *lattice*, SciPost Phys. **14**, 139 (2023), doi:[10.21468/SciPostPhys.14.6.139](https://doi.org/10.21468/SciPostPhys.14.6.139).
- 1167 [49] E. L. Weerda and M. Rizzi, *Fractional quantum hall states with variational projected*
1168 *entangled-pair states: a study of the bosonic harper-hofstadter model* (2023), [https://](https://arxiv.org/abs/2309.12811)
1169 arxiv.org/abs/2309.12811.
- 1170 [50] J. Hasik and G. B. Mbeng, *peps-torch: A differentiable tensor network library for two-*
1171 *dimensional lattice models*, GitHub, <https://github.com/jurajHasik/peps-torch>.
- 1172 [51] B. Ponsioen, *AD-PEPS*, GitHub, <https://github.com/b1592/ad-peps>.
- 1173 [52] J. Gray, *quimb: A python package for quantum information and many-body calculations*,
1174 Journal of Open Source Software **3**(29), 819, doi:[10.21105/joss.00819](https://doi.org/10.21105/joss.00819).
- 1175 [53] PEPSKit contributors, *PEPSKit: Tools for working with projected entangled-pair states.*,
1176 GitHub, <https://github.com/quantumghent/PEPSKit.jl>.
- 1177 [54] J. Dubail and N. Read, *Tensor network trial states for chiral topological phases in two*
1178 *dimensions and a no-go theorem in any dimension*, Phys. Rev. B **92**, 205307 (2015),
1179 doi:[10.1103/PhysRevB.92.205307](https://doi.org/10.1103/PhysRevB.92.205307).
- 1180 [55] Y. Xu, S. Capponi, J.-Y. Chen, L. Vanderstraeten, J. Hasik, A. H. Nevidomskyy, M. Mam-
1181 brini, K. Penc and D. Poilblanc, *Phase diagram of the chiral su(3) antiferromagnet on the*
1182 *kagome lattice*, Phys. Rev. B **108**, 195153 (2023), doi:[10.1103/PhysRevB.108.195153](https://doi.org/10.1103/PhysRevB.108.195153).
- 1183 [56] J. Naumann and E. L. Weerda, *variPEPS – a versatile tensor network library for varia-*
1184 *tional ground state simulations in two spatial dimensions*, GitHub, [https://github.com/](https://github.com/variPEPS)
1185 [variPEPS](https://github.com/variPEPS).
- 1186 [57] J. Naumann, P. Schmoll, F. Wilde and F. Krein, *variPEPS (Python version)*, Zenodo,
1187 doi:[10.5281/ZENODO.10852390](https://doi.org/10.5281/ZENODO.10852390).
- 1188 [58] E. L. Weerda, *VariPEPS.jl*, Zenodo, doi:[10.5281/zenodo.10974459](https://doi.org/10.5281/zenodo.10974459).
- 1189 [59] J. Hasik and P. Corboz, *Incommensurate order with translationally invariant projected*
1190 *entangled-pair states: Spiral states and quantum spin liquid on the anisotropic triangular*
1191 *lattice* (2023), <https://arxiv.org/abs/2311.05534>.
- 1192 [60] P. Corboz, J. Jordan and G. Vidal, *Simulation of fermionic lattice models in two dimensions*
1193 *with projected entangled-pair states: Next-nearest neighbor hamiltonians*, Phys. Rev. B **82**,
1194 245119 (2010), doi:[10.1103/PhysRevB.82.245119](https://doi.org/10.1103/PhysRevB.82.245119).
- 1195 [61] P. Corboz, T. M. Rice and M. Troyer, *Competing states in the t-j model: Uni-*
1196 *form d-wave state versus stripe state*, Phys. Rev. Lett. **113**, 046402 (2014),
1197 doi:[10.1103/PhysRevLett.113.046402](https://doi.org/10.1103/PhysRevLett.113.046402).
- 1198 [62] M. T. Fishman, L. Vanderstraeten, V. Zauner-Stauber, J. Haegeman and F. Verstraete,
1199 *Faster methods for contracting infinite two-dimensional tensor networks*, Phys. Rev. B **98**,
1200 235148 (2018), doi:[10.1103/PhysRevB.98.235148](https://doi.org/10.1103/PhysRevB.98.235148).

- 1201 [63] W. Lan and G. Evenbly, *Reduced contraction costs of corner-transfer methods for PEPS*
1202 (2023), <https://arxiv.org/abs/2306.08212>.
- 1203 [64] A. Francuz, N. Schuch and B. Vanhecke, *Stable and efficient differentiation of tensor*
1204 *network algorithms* (2023), <https://arxiv.org/abs/2311.11894>.
- 1205 [65] J. Nocedal and S. J. Wright, *Calculating Derivatives*, pp. 193–219, Springer New York,
1206 New York, NY, ISBN 978-0-387-40065-5, doi:[10.1007/978-0-387-40065-5_8](https://doi.org/10.1007/978-0-387-40065-5_8) (2006).
- 1207 [66] M. Hestenes and E. Stiefel, *Methods of conjugate gradients for solving linear systems*, J.
1208 Res. Natl. Bur. Stand. **49**(6), 409 (1952), doi:[10.6028/jres.049.044](https://doi.org/10.6028/jres.049.044).
- 1209 [67] R. Fletcher, *Function minimization by conjugate gradients*, Comp. J. **7**(2), 149 (1964),
1210 doi:[10.1093/comjnl/7.2.149](https://doi.org/10.1093/comjnl/7.2.149).
- 1211 [68] Y. H. Dai and Y. Yuan, *A nonlinear conjugate gradient method with a strong global conver-*
1212 *gence property*, SIAM J. Opt. **10**(1), 177 (1999), doi:[10.1137/s1052623497318992](https://doi.org/10.1137/s1052623497318992).
- 1213 [69] E. Polak and G. Ribiere, *Note sur la convergence de méthodes de directions conjuguées*,
1214 *Revue française d'informatique et de recherche opérationnelle. Série rouge* **3**(16), 35
1215 (1969), doi:[10.1051/m2an/196903r100351](https://doi.org/10.1051/m2an/196903r100351).
- 1216 [70] W. W. Hager and H. Zhang, *A new conjugate gradient method with guaranteed descent*
1217 *and an efficient line search*, SIAM J. Opt. **16**(1), 170 (2005), doi:[10.1137/030601880](https://doi.org/10.1137/030601880).
- 1218 [71] H. Matthies and G. Strang, *The solution of nonlinear finite element equations*, Int. J.
1219 Num. Meth. Eng. **14**(11), 1613 (1979), doi:[10.1002/nme.1620141104](https://doi.org/10.1002/nme.1620141104).
- 1220 [72] J. Nocedal, *Updating quasi-newton matrices with limited storage*, Math. Comp. **35**(151),
1221 773 (1980), doi:[10.1090/s0025-5718-1980-0572855-7](https://doi.org/10.1090/s0025-5718-1980-0572855-7).
- 1222 [73] C. G. Broyden, *The convergence of a class of double-rank minimization algorithms 1. gen-*
1223 *eral considerations*, IMA J. Appl. Math. **6**(1), 76 (1970), doi:[10.1093/imamat/6.1.76](https://doi.org/10.1093/imamat/6.1.76).
- 1224 [74] R. Fletcher, *A new approach to variable metric algorithms*, Comp. J. **13**(3), 317 (1970),
1225 doi:[10.1093/comjnl/13.3.317](https://doi.org/10.1093/comjnl/13.3.317).
- 1226 [75] D. Goldfarb, *A family of variable-metric methods derived by variational means*, Math.
1227 Comp. **24**(109), 23 (1970), doi:[10.1090/s0025-5718-1970-0258249-6](https://doi.org/10.1090/s0025-5718-1970-0258249-6).
- 1228 [76] D. F. Shanno, *Conditioning of quasi-newton methods for function minimization*, Math.
1229 Comp. **24**(111), 647 (1970), doi:[10.1090/s0025-5718-1970-0274029-x](https://doi.org/10.1090/s0025-5718-1970-0274029-x).
- 1230 [77] J. Nocedal and S. J. Wright, *Fundamentals of Unconstrained Optimization*, pp. 10–29,
1231 Springer New York, New York, NY, ISBN 978-0-387-40065-5, doi:[10.1007/978-0-387-](https://doi.org/10.1007/978-0-387-40065-5_2)
1232 [40065-5_2](https://doi.org/10.1007/978-0-387-40065-5_2) (2006).
- 1233 [78] D. C. Liu and J. Nocedal, *On the limited memory bfgs method for large scale optimization*,
1234 Math. Prog. **45**(1-3), 503 (1989), doi:[10.1007/bf01589116](https://doi.org/10.1007/bf01589116).
- 1235 [79] L. Armijo, *Minimization of Functions Having Lipschitz Continuous First Partial Deriva-*
1236 *tives*, Pac. J. Math. **16**(1), 1 (1966), doi:[10.2140/pjm.1966.16.1](https://doi.org/10.2140/pjm.1966.16.1).
- 1237 [80] P. Wolfe, *Convergence conditions for ascent methods*, SIAM Rev. **11**(2), 226 (1969),
1238 doi:[10.1137/1011036](https://doi.org/10.1137/1011036).

- 1239 [81] P. Wolfe, *Convergence conditions for ascent methods. ii: Some corrections*, SIAM Rev.
1240 **13**(2), 185 (1971), doi:[10.1137/1013035](https://doi.org/10.1137/1013035).
- 1241 [82] J. Nocedal and S. J. Wright, *Line Search Methods*, pp. 30–65, Springer New York, New
1242 York, NY, ISBN 978-0-387-40065-5, doi:[10.1007/978-0-387-40065-5_3](https://doi.org/10.1007/978-0-387-40065-5_3) (2006).
- 1243 [83] J. Haegeman, *KrylovKit*, Zenodo, doi:[10.5281/zenodo.10884302](https://doi.org/10.5281/zenodo.10884302) (2024).
- 1244 [84] IterativeSolvers contributors, *IterativeSolvers: A Julia package that provides iterative
1245 algorithms for solving linear systems, eigensystems, and singular value problems.*, GitHub,
1246 <https://github.com/JuliaLinearAlgebra/IterativeSolvers.jl>.
- 1247 [85] U. Schollwöck, *The density-matrix renormalization group in the age of matrix product
1248 states*, Annals of Physics **326**(1), 96 (2011).
- 1249 [86] H. C. Jiang, Z. Y. Weng and T. Xiang, *Accurate determination of tensor network state
1250 of quantum lattice models in two dimensions*, Phys. Rev. Lett. **101**, 090603 (2008),
1251 doi:[10.1103/PhysRevLett.101.090603](https://doi.org/10.1103/PhysRevLett.101.090603).
- 1252 [87] J. Liu, F. Wilde, A. A. Mele, L. Jiang and J. Eisert, *Stochastic noise can be helpful for
1253 variational quantum algorithms* (2023), <https://arxiv.org/abs/2210.06723>.
- 1254 [88] M. Rader and A. M. Läuchli, *Finite Correlation Length Scaling in Lorentz-
1255 Invariant Gapless iPEPS Wave Functions*, Phys. Rev. X **8**, 031030 (2018),
1256 doi:[10.1103/PhysRevX.8.031030](https://doi.org/10.1103/PhysRevX.8.031030).
- 1257 [89] P. Corboz, P. Czarnik, G. Kapteijns and L. Tagliacozzo, *Finite correlation length scal-
1258 ing with infinite projected entangled-pair states*, Phys. Rev. X **8**, 031031 (2018),
1259 doi:[10.1103/PhysRevX.8.031031](https://doi.org/10.1103/PhysRevX.8.031031).
- 1260 [90] P. Czarnik and P. Corboz, *Finite correlation length scaling with infinite projected
1261 entangled pair states at finite temperature*, Phys. Rev. B **99**, 245107 (2019),
1262 doi:[10.1103/PhysRevB.99.245107](https://doi.org/10.1103/PhysRevB.99.245107).
- 1263 [91] B. Vanhecke, J. Hasik, F. Verstraete and L. Vanderstraeten, *Scaling hypothe-
1264 sis for projected entangled-pair states*, Phys. Rev. Lett. **129**, 200601 (2022),
1265 doi:[10.1103/PhysRevLett.129.200601](https://doi.org/10.1103/PhysRevLett.129.200601).
- 1266 [92] B. Ponsioen, J. Hasik and P. Corboz, *Improved summations of n-point correlation functions
1267 of projected entangled-pair states* (2023), <https://arxiv.org/abs/2306.13327>.
- 1268 [93] I. V. Lukin and A. G. Sotnikov, *Corner transfer matrix renormalization group approach in
1269 the zoo of archimedean lattices* (2024), <https://arxiv.org/abs/2401.07274>.
- 1270 [94] A. Kitaev, *Anyons in an exactly solved model and beyond*, Ann. Phys. **321**(1), 2 (2006),
1271 doi:[10.1016/j.aop.2005.10.005](https://doi.org/10.1016/j.aop.2005.10.005).
- 1272 [95] Z. Y. Xie, J. Chen, J. F. Yu, X. Kong, B. Normand and T. Xiang, *Tensor renormalization
1273 of quantum many-body systems using projected entangled simplex states*, Phys. Rev. X **4**,
1274 011025 (2014), doi:[10.1103/PhysRevX.4.011025](https://doi.org/10.1103/PhysRevX.4.011025).
- 1275 [96] P. Schmoll, S. S. Jahromi, M. Hörmann, M. Mühlhauser, K. P. Schmidt and R. Orús,
1276 *Fine grained tensor network methods*, Phys. Rev. Lett. **124**, 200603 (2020),
1277 doi:[10.1103/PhysRevLett.124.200603](https://doi.org/10.1103/PhysRevLett.124.200603).
- 1278 [97] P. W. Anderson, *Limits on the energy of the antiferromagnetic ground state*, Phys. Rev.
1279 **83**, 1260 (1951), doi:[10.1103/PhysRev.83.1260](https://doi.org/10.1103/PhysRev.83.1260).

- 1280 [98] I. V. Lukin and A. G. Sotnikov, *Variational optimization of tensor-network states with*
1281 *the honeycomb-lattice corner transfer matrix*, Phys. Rev. B **107**(5), 054424 (2023),
1282 doi:[10.1103/physrevb.107.054424](https://doi.org/10.1103/physrevb.107.054424).
- 1283 [99] D. J. J. Farnell, O. Götze, J. Richter, R. F. Bishop and P. H. Y. Li, *Quantum*
1284 *spin- $\frac{1}{2}$ antiferromagnets on Archimedean lattices: The route from semiclassical mag-*
1285 *netic order to nonmagnetic quantum states*, Phys. Rev. B **89**(18), 184407 (2014),
1286 doi:[10.1103/physrevb.89.184407](https://doi.org/10.1103/physrevb.89.184407).
- 1287 [100] F. J. Jiang, *High precision determination of the low-energy constants for the two-*
1288 *dimensional quantum heisenberg model on the honeycomb lattice*, Europ. Phys. J. B
1289 **85**(12), 402 (2012), doi:[10.1140/epjb/e2012-30784-7](https://doi.org/10.1140/epjb/e2012-30784-7).
- 1290 [101] R. Ganesh, J. van den Brink and S. Nishimoto, *Deconfined criticality in the frus-*
1291 *trated heisenberg honeycomb antiferromagnet*, Phys. Rev. Lett. **110**, 127203 (2013),
1292 doi:[10.1103/PhysRevLett.110.127203](https://doi.org/10.1103/PhysRevLett.110.127203).
- 1293 [102] S.-S. Gong, D. N. Sheng, O. I. Motrunich and M. P. A. Fisher, *Phase diagram of the*
1294 *spin- $\frac{1}{2}$ J_1 - J_2 Heisenberg model on a honeycomb lattice*, Phys. Rev. B **88**, 165138 (2013),
1295 doi:[10.1103/PhysRevB.88.165138](https://doi.org/10.1103/PhysRevB.88.165138).
- 1296 [103] H. J. Liao, Z. Y. Xie, J. Chen, Z. Y. Liu, H. D. Xie, R. Z. Huang, B. Normand and T. Xiang,
1297 *Gapless spin-liquid ground state in the $s = 1/2$ kagome antiferromagnet*, Phys. Rev. Lett.
1298 **118**, 137202 (2017), doi:[10.1103/PhysRevLett.118.137202](https://doi.org/10.1103/PhysRevLett.118.137202).
- 1299 [104] A. M. Läuchli, J. Sudan and R. Moessner, *Spin- $\frac{1}{2}$ kagome Heisenberg antiferromagnet*
1300 *revisited*, Phys. Rev. B **100**(15), 155142 (2019), doi:[10.1103/physrevb.100.155142](https://doi.org/10.1103/physrevb.100.155142).
- 1301 [105] S. Yan, D. A. Huse and S. R. White, *Spin-Liquid Ground State of the S*
1302 *$= 1/2$ Kagome Heisenberg Antiferromagnet*, Science **332**(6034), 1173 (2011),
1303 doi:[10.1126/science.1201080](https://doi.org/10.1126/science.1201080).
- 1304 [106] S. Depenbrock, I. P. McCulloch and U. Schollwöck, *Nature of the spin-liquid ground state*
1305 *of the $s = 1/2$ heisenberg model on the kagome lattice*, Phys. Rev. Lett. **109**, 067201
1306 (2012), doi:[10.1103/PhysRevLett.109.067201](https://doi.org/10.1103/PhysRevLett.109.067201).
- 1307 [107] Y.-C. He, M. P. Zaletel, M. Oshikawa and F. Pollmann, *Signatures of Dirac Cones in*
1308 *a DMRG Study of the Kagome Heisenberg Model*, Phys. Rev. X **7**, 031020 (2017),
1309 doi:[10.1103/PhysRevX.7.031020](https://doi.org/10.1103/PhysRevX.7.031020).
- 1310 [108] N. Astrakhantsev, F. Ferrari, N. Niggemann, T. Müller, A. Chauhan, A. Kshetrimayum,
1311 P. Ghosh, N. Regnault, R. Thomale, J. Reuther, T. Neupert and Y. Iqbal, *Pinwheel valence*
1312 *bond crystal ground state of the spin- $\frac{1}{2}$ heisenberg antiferromagnet on the shuriken lattice*,
1313 Phys. Rev. B **104**, L220408 (2021), doi:[10.1103/PhysRevB.104.L220408](https://doi.org/10.1103/PhysRevB.104.L220408).
- 1314 [109] P. Schmoll, A. Kshetrimayum, J. Naumann, J. Eisert and Y. Iqbal, *Tensor network study of*
1315 *the spin- $\frac{1}{2}$ Heisenberg antiferromagnet on the shuriken lattice*, Phys. Rev. B **107**, 064406
1316 (2023), doi:[10.1103/PhysRevB.107.064406](https://doi.org/10.1103/PhysRevB.107.064406).
- 1317 [110] M. Fujihala, K. Morita, R. Mole, S. Mitsuda, T. Tohyama, S.-I. Yano, D. Yu, S. Sota,
1318 T. Kuwai, A. Koda, H. Okabe, H. Lee *et al.*, *Gapless spin liquid in a square-kagome lattice*
1319 *antiferromagnet*, Nature Comm. **11**, 3429 (2020), doi:[10.1038/s41467-020-17235-z](https://doi.org/10.1038/s41467-020-17235-z).
- 1320 [111] Q. Li, H. Li, J. Zhao, H.-G. Luo and Z. Y. Xie, *Magnetization of the spin- $\frac{1}{2}$ heisen-*
1321 *berg antiferromagnet on the triangular lattice*, Phys. Rev. B **105**, 184418 (2022),
1322 doi:[10.1103/PhysRevB.105.184418](https://doi.org/10.1103/PhysRevB.105.184418).

- 1323 [112] D. A. Huse and V. Elser, *Simple variational wave functions for two-dimensional*
1324 *heisenberg spin- $1/2$ antiferromagnets*, Phys. Rev. Lett. **60**, 2531 (1988),
1325 doi:10.1103/PhysRevLett.60.2531.
- 1326 [113] S. R. White and A. L. Chernyshev, *Neél order in square and triangular lattice heisenberg*
1327 *models*, Phys. Rev. Lett. **99**, 127004 (2007), doi:10.1103/PhysRevLett.99.127004.
- 1328 [114] S. Yunoki and S. Sorella, *Two spin liquid phases in the spatially anisotropic triangular*
1329 *heisenberg model*, Phys. Rev. B **74**, 014408 (2006), doi:10.1103/PhysRevB.74.014408.
- 1330 [115] V. Zauner, D. Draxler, L. Vanderstraeten, M. Degroote, J. Haegeman, M. M. Rams, V. Sto-
1331 jevic, N. Schuch and F. Verstraete, *Transfer matrices and excitations with matrix product*
1332 *states*, New J. Phys. **17**, 053002 (2015), doi:10.1088/1367-2630/17/5/053002.
- 1333 [116] L. Vanderstraeten, M. Mariën, F. Verstraete and J. Haegeman, *Excitations and the*
1334 *tangent space of projected entangled-pair states*, Phys. Rev. B **92**, 201111 (2015),
1335 doi:10.1103/PhysRevB.92.201111.
- 1336 [117] L. Vanderstraeten, J. Haegeman and F. Verstraete, *Simulating excitation spec-*
1337 *tra with projected entangled-pair states*, Phys. Rev. B **99**, 165121 (2019),
1338 doi:10.1103/PhysRevB.99.165121.
- 1339 [118] B. Ponsioen and P. Corboz, *Excitations with projected entangled pair states us-*
1340 *ing the corner transfer matrix method*, Phys. Rev. B **101**, 195109 (2020),
1341 doi:10.1103/PhysRevB.101.195109.
- 1342 [119] W.-L. Tu, L. Vanderstraeten, N. Schuch, H.-Y. Lee, N. Kawashima and J.-Y. Chen, *Gen-*
1343 *erating function for projected entangled-pair states* (2023), [https://arxiv.org/abs/2307.](https://arxiv.org/abs/2307.08083)
1344 [08083](https://arxiv.org/abs/2307.08083).
- 1345 [120] T. Barthel, C. Pineda and J. Eisert, *Contraction of fermionic operator circuits and*
1346 *the simulation of strongly correlated fermions*, Phys. Rev. A **80**, 042333 (2009),
1347 doi:10.1103/PhysRevA.80.042333.
- 1348 [121] P. Corboz, G. Evenbly, F. Verstraete and G. Vidal, *Simulation of interacting*
1349 *fermions with entanglement renormalization*, Phys. Rev. A **81**, 010303 (2010),
1350 doi:10.1103/PhysRevA.81.010303.
- 1351 [122] C. Wille, O. Buerschaper and J. Eisert, *Fermionic topological quantum states as tensor*
1352 *networks*, Phys. Rev. B **95**, 245127 (2017), doi:10.1103/PhysRevB.95.245127.
- 1353 [123] N. Bultinck, D. J. Williamson, J. Haegeman and F. Verstraete, *Fermionic pro-*
1354 *jected entangled-pair states and topological phases*, J. Phys. A **51**, 025202 (2017),
1355 doi:10.1088/1751-8121/aa99cc.
- 1356 [124] P. Corboz, R. Orús, B. Bauer and G. Vidal, *Simulation of strongly correlated fermions in*
1357 *two spatial dimensions with fermionic projected entangled-pair states*, Phys. Rev. B **81**,
1358 165104 (2010), doi:10.1103/PhysRevB.81.165104.
- 1359 [125] I. Pižorn and F. Verstraete, *Fermionic implementation of projected entangled pair states*
1360 *algorithm*, Phys. Rev. B **81**, 245110 (2010), doi:10.1103/PhysRevB.81.245110.
- 1361 [126] R. Sweke, P. Boes, N. Ng, C. Sparaciari, J. Eisert and M. Goihl, *Transparent Reporting*
1362 *of Research-Related Greenhouse Gas Emissions Through the Scientific CO₂nduct Initiative*,
1363 Comm. Phys. **5**(1), 150 (2022), doi:10.1038/s42005-022-00930-2.

- 1364 [127] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula,
1365 A. Paszke, J. VanderPlas, S. Wanderman-Milne and Q. Zhang, *JAX: composable trans-*
1366 *formations of Python+NumPy programs*, GitHub, <http://github.com/google/jax>.
- 1367 [128] R. Frostig, M. Johnson and C. Leary, *Compiling machine learning programs via high-level*
1368 *tracing* (2018), <https://mlsys.org/Conferences/2019/doc/2018/146.pdf>.
- 1369 [129] Jutho, Lukas, M. Hauru, ho oto, maartenvd, Gertian, J. TagBot, S. Carlström and Xi-
1370 aoyu, *Jutho/TensorKit.jl: v0.12.2*, Zenodo, doi:[10.5281/zenodo.10574897](https://doi.org/10.5281/zenodo.10574897) (2024).
- 1371 [130] Quantum Group UGent, *Open-source software and notebooks*, <https://quantumghent.github.io/software/>.
- 1373 [131] Zygote contributors, *Zygote: Source-to-source automatic differentiation (AD) in Julia,*
1374 *and the next-gen AD system for the Flux differentiable programming framework.*, GitHub,
1375 <https://github.com/FluxML/Zygote.jl>.
- 1376 [132] J. S. Centre, *Jureca: Data centric and booster modules implementing the modular super-*
1377 *computing architecture at jülich supercomputing centre*, *J. Large-Scale Res. Fac.* **7**, A182
1378 (2021), doi:[10.17815/jlsrf-7-182](https://doi.org/10.17815/jlsrf-7-182).
- 1379 [133] M. B. Giles, *An extended collection of matrix derivative results for forward and reverse*
1380 *mode algorithmic differentiation* (2008), [https://people.maths.ox.ac.uk/gilesm/files/](https://people.maths.ox.ac.uk/gilesm/files/NA-08-01.pdf)
1381 [NA-08-01.pdf](https://people.maths.ox.ac.uk/gilesm/files/NA-08-01.pdf).
- 1382 [134] H. Xie, J.-G. Liu and L. Wang, *Automatic differentiation of dominant eigen-*
1383 *solver and its applications in quantum physics*, *Phys. Rev. B* **101**, 245139 (2020),
1384 doi:[10.1103/physrevb.101.245139](https://doi.org/10.1103/physrevb.101.245139).
- 1385 [135] Z.-Q. Wan and S.-X. Zhang, *Automatic differentiation for complex valued SVD* (2019),
1386 <https://arxiv.org/abs/1909.02659>.
- 1387 [136] B. Christianson, *Reverse accumulation and attractive fixed points*, *Opt. Meth. Soft.* **3**(4),
1388 311 (1994), doi:[10.1080/10556789408805572](https://doi.org/10.1080/10556789408805572).
- 1389 [137] Z. Kolter, D. Duvenaud and M. Johnson, *Deep implicit layers - Neural ODEs, deep*
1390 *equilibrium models, and beyond* (2020), [http://implicit-layers-tutorial.org/implicit_](http://implicit-layers-tutorial.org/implicit_functions/)
1391 [functions/](http://implicit-layers-tutorial.org/implicit_functions/).
- 1392 [138] M. J. Johnson and The JAX Authors, *Custom derivative rules for JAX-transformable*
1393 *Python functions* (2020), [https://jax.readthedocs.io/en/latest/notebooks/Custom_](https://jax.readthedocs.io/en/latest/notebooks/Custom_derivative_rules_for_Python_code.html)
1394 [derivative_rules_for_Python_code.html](https://jax.readthedocs.io/en/latest/notebooks/Custom_derivative_rules_for_Python_code.html).