

PAPER • OPEN ACCESS

Symmetric tensor networks for generative modeling and constrained combinatorial optimization

To cite this article: Javier Lopez-Piqueres *et al* 2023 *Mach. Learn.: Sci. Technol.* 4 035009

View the [article online](#) for updates and enhancements.

You may also like

- [Some methods for identifying redundant constraints in linear programming](#)
Y Estinmgsih, Farikhin and R H Tjahjana
- [Investigating the Hubble Constant Tension: Two Numbers in the Standard Cosmological Model](#)
Weikang Lin, Katherine J. Mack and Liqiang Hou
- [An Efficient Local Search algorithm for Over-Constrained CSPs](#)
Jufang Li and Yan Chi



PAPER

OPEN ACCESS

RECEIVED
1 April 2023REVISED
25 May 2023ACCEPTED FOR PUBLICATION
19 June 2023PUBLISHED
21 July 2023

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Symmetric tensor networks for generative modeling and constrained combinatorial optimization

Javier Lopez-Piqueres^{1,2} , Jing Chen^{1,*} and Alejandro Perdomo-Ortiz^{3,*} ¹ Zapata Computing Inc., 100 Federal Street, Boston, MA 02110, United States of America² Department of Physics, University of Massachusetts, Amherst, MA 01003, United States of America³ Zapata Computing Canada Inc., 25 Adelaide St East, M5C 3A1 Toronto, ON, Canada

* Authors to whom any correspondence should be addressed.

E-mail: jing.chen@zapatacomputing.com and alejandro@zapatacomputing.com**Keywords:** generative learning, combinatorial optimization, tensor network, constraints, $U(1)$ symmetry

Abstract

Constrained combinatorial optimization problems abound in industry, from portfolio optimization to logistics. One of the major roadblocks in solving these problems is the presence of non-trivial hard constraints which limit the valid search space. In some heuristic solvers, these are typically addressed by introducing certain Lagrange multipliers in the cost function, by relaxing them in some way, or worse yet, by generating many samples and only keeping valid ones, which leads to very expensive and inefficient searches. In this work, we encode arbitrary integer-valued equality constraints of the form $A\vec{x} = \vec{b}$, directly into $U(1)$ symmetric tensor networks (TNs) and leverage their applicability as quantum-inspired generative models to assist in the search of solutions to combinatorial optimization problems within the *generator-enhanced optimization* framework of Alcazar *et al* (2021 arXiv:2101.06250). This allows us to exploit the generalization capabilities of TN generative models while constraining them so that they only output valid samples. Our constrained TN generative model efficiently captures the constraints by reducing number of parameters and computational costs. We find that at tasks with constraints given by arbitrary equalities, symmetric matrix product states outperform their standard unconstrained counterparts at finding novel and better solutions to combinatorial optimization problems.

1. Introduction

Tensor networks (TNs) have emerged in the past few years as a very powerful tool to address challenging problems in high-dimensional spaces, such as quantum many-body physics [1], electronic structure of chemical compounds [2], supervised learning [3], generative modeling [4], anomaly detection [5], combinatorial optimization [6–8], and privacy leaking [9], to name a few. The reason for their wide applicability lies in part for their well-understood mathematical structure and well-controlled expressibility.

Just as artificial neural networks (ANNs) are universal approximators of generic distributions [10], TNs can in principle model any (discrete) distribution with enough computational resources [11]. While ANNs are inspired by the functioning of the brain (in particular their expressibility determined by the number of activation functions, mimicking the firing of neurons in the brain), themselves nonlinear functions of their input, TNs are inspired by quantum mechanics, arguably a theory with richer mathematical structure than that of ANNs. For instance, the Hilbert space structure of quantum mechanics allows to exploit global *internal* symmetries in quantum states through the use of representation theory. This can have important consequences when we aim to efficiently express probability distributions with redundancies that appear as a result of such symmetries.

It may appear at first sight that exploiting internal symmetries such as $U(1)$ or $SU(2)$ would bear little importance in problems outside physics. In this work, we show that by exploiting $U(1)$ symmetry in TN states it is possible to encode arbitrary integer-valued equalities of the form $A\vec{x} = \vec{b}$ in a probabilistic model,

extending the family of so-called tensor network Born machine (TNBM) models [4]. Therefore, we refer to our new family of symmetric quantum-inspired TN generative models as s-TNBMs.

This type of equalities appear in many combinatorial optimization problems with *hard* constraints, such as in portfolio optimization, warehouse location, scheduling, transportation, to name a few [12, 13]. Current state-of-the-art methods at dealing with combinatorial optimization problems with hard-constraints, such as MIXED-INTEGER LINEAR PROGRAMMING (MILP) solvers [14, 15], simulated annealing solvers [16], or more recently, graph neural networks (GNNs) [17–19] and fully quantum models such as the Quantum Approximate Optimization Algorithm (QAOA) [20], suffer from at least one of the following two drawbacks: they must explicitly break the integrality constraint to then project back to the nearest integer solution (as done in standard MILP solvers relying on the *branch-and-bound* method [13]), and/or they only work for cost functions expressed as polynomials in the binary variables, with linear or quadratic being the most common ones. Here, we focus on a recent proposal, the generator-enhanced optimization (GEO) framework [21], which bypasses both of these limitations.

GEO solvers leverage classical or quantum generative models to assist the solution of combinatorial optimization problems, by learning the correlations from observations (i.e. bitstrings and their corresponding costs) and proposing new bitstring solution candidates which resemble those in the lowest cost sector. This is done by reweighting the seen information, which is set as the dataset to train the generative model. This reweighting helps to give more importance to those samples in the training dataset with the lowest cost (see section 2 for more details). It was shown in [21] the power of quantum-inspired generative models based on TNBMs within GEO, matching the performance of state-of-the-art metaheuristics tailored in the last 30 years for a specific variant of the portfolio optimization problem. These remarkable results were obtained without imposing any bias in the TNBM model to enforce the cardinality constraints (i.e. fixed desired number k of assets in the portfolio out of N) of the combinatorial problems.

In this work, we exploit the rather simple structure of TNs, in particular their linearity and local connectivity, to encode arbitrary equality constraints, cardinality being a specific realization of these, for the purpose of both generative modeling and combinatorial optimization. We do so by mapping the problem of finding solutions satisfying the constraints $A\vec{x} = \vec{b}$ to that of finding the set of *quantum numbers* (QNs) appearing in symmetric TNs. In this context, QNs effectively correspond to labels of different irreducible representations, in this case of $U(1)$, defined on each local Hilbert space where each tensor is supported.

For equality constraints with high-degree of symmetry, where the coefficients in matrix A have low variance, we find that symmetric TNs correspond to an efficient encoding of the valid space. For a subset of these, such as cardinality type constraints appearing in problems such as portfolio optimization [21], one can write down explicitly the TN ansatz for the entire valid space, which can be used as an initial ansatz for finding optimal solutions. For arbitrary integer-valued equalities finding the solution space is likely #P-complete (since its associated decision problem, (0,1)-INTEGER PROGRAMMING, is already NP-complete [22]), but we provide a novel message-passing algorithm with QNs acting as *messages*, that is able to encode a large fraction of these solutions within a few QNs, thereby allowing to generate many unseen solutions satisfying $A\vec{x} = \vec{b}$ when given access to some (valid) training data, even when the entries of A and \vec{b} are random integers. At the same time, by exploiting the *dimensionality* of each QN, the resulting TN ansatz is able to capture not only the constraints efficiently, but also biases in the dataset, therefore favoring certain solutions over others.

Moreover, by constraining the TN so as to only sample within the valid subspace, we guarantee computational savings, in both memory usage by exploiting block-sparsity of the tensors as well as computational time by guaranteeing that only the needed tensor components get updated. Beyond the ability to impose arbitrary equality constraints, TNs enjoy sophisticated optimization algorithms, such as the density matrix renormalization group (DMRG) algorithm [23, 24] for matrix product states (MPSs), as well as perfect sampling [25] (which avoid long autocorrelation times typical of Monte Carlo based sampling algorithms).

While we envision other uses of constrained TNs as efficient representations of constrained datasets in various contexts, we find that some of its most crucial advantages are captured within the GEO framework of [21]. This is because GEO exploits the presence of existing training data to generate new and better solutions to combinatorial optimization problems by capturing biases in the training set (corresponding to bitstrings of different costs), which our constrained TN approach is capable of incorporating when given data subject to equality constraints.

The structure of the paper is as follows. In section 2. we state the problem we are interested in, namely combinatorial optimization problems subject to equality constraints, and describe our approach to the problem using a generative model based on TNs. In section 3 we give a brief introduction to symmetric TNs placing emphasis on symmetric MPS, a special TN structure which will be used in our numerical experiments. In section 4 we explain the main encoding and generative algorithms. In section 5 we present

various results showcasing the types of canonical problems symmetric MPS can be applied to and compare its performance against vanilla versions (i.e. without imposing any constraint on them). In section 6 we give strong arguments in favor of using TNs for describing symmetric states over ANN architectures, implying in particular the efficiency of symmetric TNs at capturing arbitrary equality constraints when restricted to $U(1)$ symmetric states. Section 7 presents our conclusion, and section 8 gives an outlook for future work.

2. Problem statement

In combinatorial optimization problems, one is interested in finding the minimum of a cost function over a specific set of instances, called valid (or *feasible*) space. Often the set of valid instances grows exponentially with the problem size (such as in the presence of cardinality constraints), which prevents listing exhaustively all possible candidates. Although not all combinatorial problems have explicit constraints, the class of real-world problems that we focus on here are those which can be written as:

$$\begin{aligned} \min \mathcal{C}(\vec{x}), \\ \text{subject to } A\vec{x} = \vec{b}, \end{aligned} \quad (1)$$

where $\mathcal{C} : \{0, 1\}^N \rightarrow \mathbb{R}$, $\vec{x} \in \{0, 1\}^N$, and $A \in \mathbb{Z}^{m \times N}$, $\vec{b} \in \mathbb{Z}^m$. (Some of these conditions may be relaxed in principle and our construction still applies; for instance, we may convert each integer coefficient to a rational and viceversa via a rescaling of all coefficients in each equality. Similarly, we may work with general integer-valued variables instead of just binary ones by accordingly increasing the dimension of each site vector space in the TN as we will describe in the next section). We will denote by \mathcal{S} the valid or solution space, i.e. the set of all bitstrings satisfying the equality constraints. While not all hard constraints are of the form appearing in equation (1) [26] our focus is on equality type constraints because these are the ones that best exploit the TN ansatz used in this work.

When the cost function \mathcal{C} is linear, the above statement reduces to the *standard form* of (0,1)-INTEGER LINEAR PROGRAMMING (see appendix A for an introduction into this class of problems), and when \mathcal{C} is at most quadratic, the corresponding problem can be recasted to a QUADRATIC UNCONSTRAINED BINARY OPTIMIZATION (QUBO) with penalty terms enforcing the equality constraints [27]. In the GEO framework used here, the form of $\mathcal{C}(\vec{x})$ can be arbitrary, i.e. GEO is a black-box solver. While many powerful methods exist for dealing with the minimization part of the algorithm such as *branch-and-bound* based methods for linear cost functions, and *simulated annealing* for problems with arbitrary cost functions, a major roadblock in many of these solvers is dealing precisely with equality constraints (*hard* constraints) that restrict the valid space, and the standard approaches taken are to relax both the integrality and equality constraints as in branch-and-bound based solvers [22], to impose such constraints via introducing Lagrange multipliers directly into the cost function as in most QUBO solvers [27] (which for many equalities results in an overparameterized cost function), or worse yet, to generate many samples only to keep those that meet the constraints.

Inspired by current generative modeling approaches to the task of combinatorial optimization such as the ones from [21, 28, 29], we leverage the generalization capabilities of TNBMs to address the optimization task in equation (1) while constraining them so that they only output valid samples. Crucially, we do not make any assumption on the nature of the cost function (hence our approach is not limited to linear/quadratic cost functions). Our approach builds on top of the GEO approach of [21] when subject to arbitrary integer-valued equality constraints. Specifically, denoting *constrained-GEO* the constrained extension of GEO, the problem statement we are interested in can be phrased as follows

CONSTRAINED-GEO

Input: A training dataset $\mathcal{T} \subset \{0, 1\}^N$ satisfying $A\vec{x} = \vec{b}$, $\forall \vec{x} \in \mathcal{T}$, with $A \in \mathbb{Z}^{m \times N}$, $\vec{b} \in \mathbb{Z}^m$.

Task: Find new, valid samples, $\vec{x} \in \mathcal{S} \setminus \mathcal{T}$, while minimizing $\mathcal{L} = -\sum_{\vec{x} \in \mathcal{T}} p_T(\vec{x}) \log \mathbb{P}(\vec{x})$.

Here, \mathcal{L} is the training loss function in the form of the cross-entropy or negative log likelihood (NLL), $\mathbb{P}(\vec{x})$ is the *model* distribution represented as a TNBM, i.e. $\mathbb{P}(\vec{x}) = |\Psi(\vec{x})|^2 / Z$ with $\Psi(\vec{x})$ expressed as a TN and $Z = \sum_{\vec{x} \in \{0, 1\}^N} |\Psi(\vec{x})|^2$, and $p_T(\vec{x})$ is the training distribution given by the *softmax*

$$p_T(\vec{x}) = \frac{1}{M} e^{-\mathcal{C}(\vec{x})/T}, \quad \vec{x} \in \mathcal{T}, \quad (2)$$

where $M = \sum_{\vec{x} \in \mathcal{T}} e^{-\mathcal{C}(\vec{x})/T}$ is a normalization factor, and the cost function $\mathcal{C}(\vec{x})$ is in principle arbitrary (not restricted to *polynomial* cost functions) and T is a Lagrange multiplier favoring lower cost samples as $T \rightarrow 0$. The motivation for the use of symmetric MPS (and in general, symmetric TNs) is that, as we shall show next,

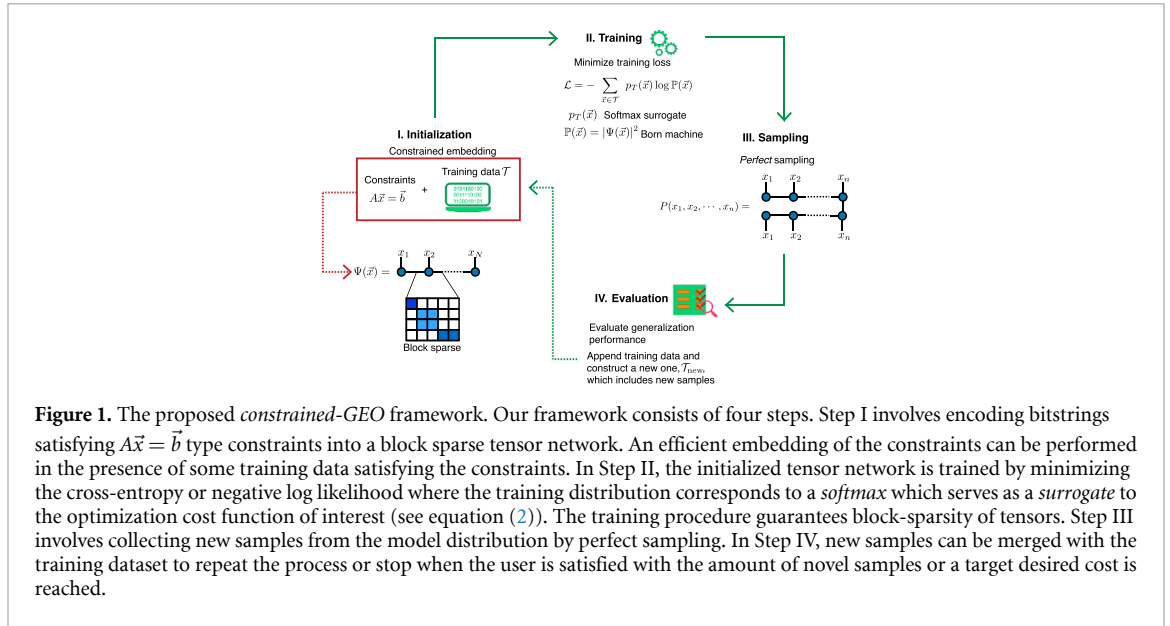


Figure 1. The proposed *constrained-GEO* framework. Our framework consists of four steps. Step I involves encoding bitstrings satisfying $A\vec{x} = \vec{b}$ type constraints into a block sparse tensor network. An efficient embedding of the constraints can be performed in the presence of some training data satisfying the constraints. In Step II, the initialized tensor network is trained by minimizing the cross-entropy or negative log likelihood where the training distribution corresponds to a *softmax* which serves as a *surrogate* to the optimization cost function of interest (see equation (2)). The training procedure guarantees block-sparsity of tensors. Step III involves collecting new samples from the model distribution by perfect sampling. In Step IV, new samples can be merged with the training dataset to repeat the process or stop when the user is satisfied with the amount of novel samples or a target desired cost is reached.

they permit to incorporate equality constraints *directly* into a generative model given in terms of a symmetric MPS, where tensors acquire a block sparse form. In figure 1 we schematically describe our proposed constrained-GEO framework for solving arbitrary combinatorial optimization problems subject to equality constraints which we will describe in more depth in section 4.

3. Symmetric TNs

In this section we give a brief introduction to $U(1)$ symmetric MPS (for more details we refer to the appendix B). Our motivation here is to set up the notation and terminology that will be used in the following sections as well as to showcase how symmetric MPS allows to encode arbitrary integer-valued equality constraints.

3.1. Vanilla and $U(1)$ symmetric MPSs

TNs are a special ansatz to represent high dimensional data, such as quantum states, based on the contraction of tensors defined on networks. The tensors, whose entries are tunable parameters, are located at the sites of the network, while the links between tensors capture the correlations among different sites and represent the amount of entanglement shared among them. It has been shown that many states of physical relevance most prominently those obeying the so-called area law of entanglement [30, 31], can be efficiently represented by such ansatz. Among the different classes of TNs, MPS corresponding to one-dimensional TNs, has found the widest application, in great part due to its low connectivity as well as the presence of a *canonical form*, which permits very optimized updating and contracting algorithms [24, 32]. A perspective that is useful when introducing *symmetries* in a TN is to view each tensor as a linear map from some *input* vector space \mathbb{V}^{in} to *output* vector space \mathbb{V}^{out} . For instance, a 3-rank tensor in an MPS may be of the form $|\alpha\rangle = \sum_{a,\beta} T_{\alpha\beta}^a |a,\beta\rangle$ where $|\alpha\rangle \in \mathbb{V}^{\text{in}}$ and $|a,\beta\rangle \in \mathbb{V}^{\text{out}}$. Here we make a distinction between upper indices appearing in each tensor component, referring to them as site (or *physical*) indices (indices that are not contracted in the TN), and lower indices corresponding to link (or *virtual*) indices (indices that are contracted and that control the expressibility of the TN ansatz). The dimension of each link's vector space is determined by the bond dimension χ , controlling the amount of *entanglement* shared across nearest neighbors. Its size determines in turn the expressibility of the resulting MPS ansatz. One of the benefits of TNs is that one can reason with them solely based on their diagrammatic representation. For the tensor above, this is represented as

$$T_{\alpha\beta}^a = \alpha \rightarrow \text{---} \bullet \xrightarrow{a} \text{---} \beta$$

where the direction of the arrows reflect whether they refer to incoming/outgoing states.

TNs not only permit to construct states based on their entanglement structure, but also allow to exploit global internal symmetries such as $U(1)$ and $SU(2)$, widely present in many problems of physical relevance. This essentially constrains the form of each tensor so as to fulfill this global requirement. The simplest kind of symmetric TN is the $U(1)$ symmetric MPS. An arbitrary vector $|v\rangle \in \mathbb{V}$ can be expanded as a linear

combination of $|n, t_n\rangle \in \mathbb{V}_n$ vectors of fixed QN or *charge* n , and degeneracy label $t_n = 1, \dots, d_n$, where d_n is the degeneracy of charge n . By going to this basis of well-defined charge, a symmetric MPS can be written in terms of *constrained* tensors of the form [33, 34]

$$T_{\alpha,\beta}^a = (T_{n_\alpha, n_\beta}^{n_a})_{t_{n_\alpha}, t_{n_\beta}}^{t_{n_a}} \cdot \delta_{n+N_{in}, N_{out}}, \tag{3}$$

where

$$N_{in} = \sum_{i \in \mathcal{I}} n_i, \quad N_{out} = \sum_{i \in \mathcal{O}} n_i, \tag{4}$$

where \mathcal{I} (\mathcal{O}) denotes the set of incoming (outgoing) indices, and n denotes the *flux* of the tensor, a label carrying the charge of the tensor. When using the canonical form (as will be the case throughout in this work), a symmetric MPS is usually chosen so that the flux of the entire MPS is only carried by one of the tensors (instead of spread across various tensors), the canonical center (also known as *orthogonality center*), while the remaining ones have flux zero ($n = 0$ in equation (3)). Thus, by exploiting global $U(1)$ symmetry we see that each tensor in the MPS factorizes into the product of two tensors. The first one, $(T_{n_\alpha, n_\beta}^{n_a})_{t_{n_\alpha}, t_{n_\beta}}^{t_{n_a}}$, gives the components of tensor T w.r.t. the degeneracy vectors $\{|t_n\rangle\}$, for that reason this tensor goes by the name of *degeneracy tensor* [33]. The second tensor, $\delta_{N_{in}, N_{out}}$ depends solely on the conservation of $U(1)$ QNs, and it is referred to as *structural tensor* (for an arbitrary symmetry group \mathcal{G} the structural tensor can take a rather nontrivial form) [33]. For fixed physical QN (fixed n_a appearing in equation (3)), the resulting matrix $T_{\alpha,\beta}^a$ will be block diagonal, with blocks of order $d_{n_\alpha} \times d_{n_\beta}$. These blocks are sometimes referred to as QN blocks. In other words, charge conservation constrains the form of the tensors to become block sparse, leading to a more efficient representation of the MPS.

3.2. Constructing $U(1)$ symmetric MPS

Having discussed the formalism of $U(1)$ symmetric MPS, we now explain how to explicitly construct a symmetric MPS that only generates bitstrings fulfilling different equality constraints of the form $A\vec{x} = \vec{b}$, with $\vec{x} \in \{0, 1\}^N$, $A \in \mathbb{Z}^{m \times N}$, and $\vec{b} \in \mathbb{Z}^m$. Our construction relies on identifying the site (*physical*) charges as the entries of the A matrix together with the empty charge \emptyset (i.e. the charge 0 in the case of integer valued scalar charges), and the total charge of the MPS (i.e. the flux) as the entries of \vec{b} . The remaining task consists in determining all link charges consistent with global charge conservation. This will automatically give us the components of each structural tensor (the degeneracy tensor can always be trivially constructed by augmenting the dimension of each QN block). We remark, however, that finding the set of all possible configurations of link charges that add up to flux \vec{b} is likely #P-hard, since its associated decision problem, (MULTIPLE) SUBSET SUM, is at least NP-hard (the variant with A and \vec{b} being one dimensional and having non-negative integer entries is for instance NP-complete [35]). Therefore, in general we will not be interested in finding all such charges for a given A, \vec{b} , but rather in finding a way to incorporate them into the MPS once these are found (finding all charges for small-sized problems may be done via exhaustive search, random search or dynamic programming).

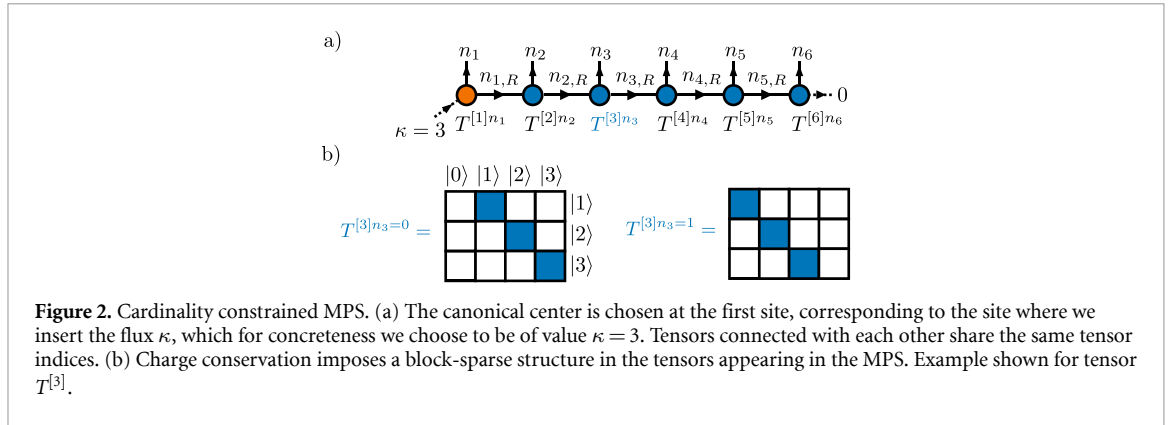
3.2.1. Cardinality constraint—an exactly solvable case

The original motivation behind the construction of symmetric MPS arose from capturing the conservation law of total spin \hat{S}_z along the z direction in spin chains (or alternatively, the conservation of particle number in systems of spinless bosons and fermions). These two types of conservation laws can be viewed as cardinality constraints of the form $\sum_i x_i = \kappa$, where $x_i \in \{0, 1\}$, and $\kappa \in \mathbb{N}_0$ s.t. $0 \leq \kappa \leq N$. We shall focus on the case $\kappa \leq N/2$ (assuming N even) since once we find a way to generate bitstrings fulfilling $\sum_i x_i = \kappa \leq N/2$ we can simply flip all bits so as to generate $\sum_i x_i = N - \kappa \geq N/2$.

Given the simplicity of the constraint, we can construct the set of all charges explicitly. For concreteness, we select the canonical center of the MPS to be located at the first site, where we insert the flux of charge κ ; see figure 2. The construction of the MPS is accomplished once we determine the charges at the link indices. To distinguish between site charges and link charges we will denote the former via n_i with $i = 1, \dots, N$, and the link charges as $n_{i,R}$, $i = 1, \dots, N - 1$ for the link charge immediately to the right of charge n_i . As mentioned at the beginning of this subsection, the site charges will be given by the coefficients of the A matrix, together with the \emptyset charge. For cardinality, $n_i \in \{\emptyset, 1\}$.

The construction of the link charges requires solving the set of equations $n_{i+1,R} = n_{i,R} - n_{i+1}$ with the boundary terms $n_{1,R} = \kappa - n_1$ and $n_{N-1,R} = n_N$. The solution is given by

$$n_{i,R} \in \begin{cases} \{\kappa, \kappa - 1, \dots, \kappa - i\}, & 1 \leq i < \kappa, \\ \{\kappa, \kappa - 1, \dots, \emptyset\}, & \kappa \leq i \leq N - \kappa, \\ \{N - i, N - i - 1, \dots, \emptyset\}, & N - \kappa < i \leq N - 1. \end{cases} \tag{5}$$



Crucially, while the solution space is given by the binomial coefficient $|S| = \binom{N}{\kappa}$ (which in particular for $\kappa \approx N/2$ and $N \gg 1$ grows as $|S| \sim \mathcal{O}(e^N/\sqrt{N})$), the number of parameters in the symmetric MPS ansatz grows as $\mathcal{O}(N(\kappa + 1))$. Thus, the symmetric MPS ansatz is an *efficient encoding* of the solution space.

3.2.2. Arbitrary single equality

Next we consider an equality of the type $\sum_i a_i x_i = b$ with $a_i, b \in \mathbb{Z}$, which appears for instance in SUBSET SUM-type problems. In this setting, we can still impose this constraint into an MPS, which is accomplished by first setting the local site charges as $n_i \in \{\emptyset, a_i\}$. As opposed to the cardinality constraint above, we cannot *a priori* determine the set of link charges for a general equality. A standard approach for this problem would be to use backtracking: construct all possible link charges solving $n_{i,R} = n_{i+1} + n_{i+1,R}$ in a forward sweep, and discard those that are not consistent with charge conservation in a backward sweep. In practice, one can do better than this and apply a *meet-in-the-middle* technique [36] whereby we locate the canonical center near the middle of the MPS and enumerate exhaustively all configurations of charges to the left and to the right, as well as their corresponding sums. Using a binary search tree on top to match configurations of charges to the left and to the right adding up to the total charge of interest brings down the time complexity from $\mathcal{O}(2^N)$ to $\mathcal{O}(N2^{N/2})$. Alternatively, when the coefficients a_i, b are non-negative one may use a standard dynamic programming approach with pseudopolynomial runtime $\mathcal{O}(Nb)$ [37].

3.2.3. Arbitrary set of equalities

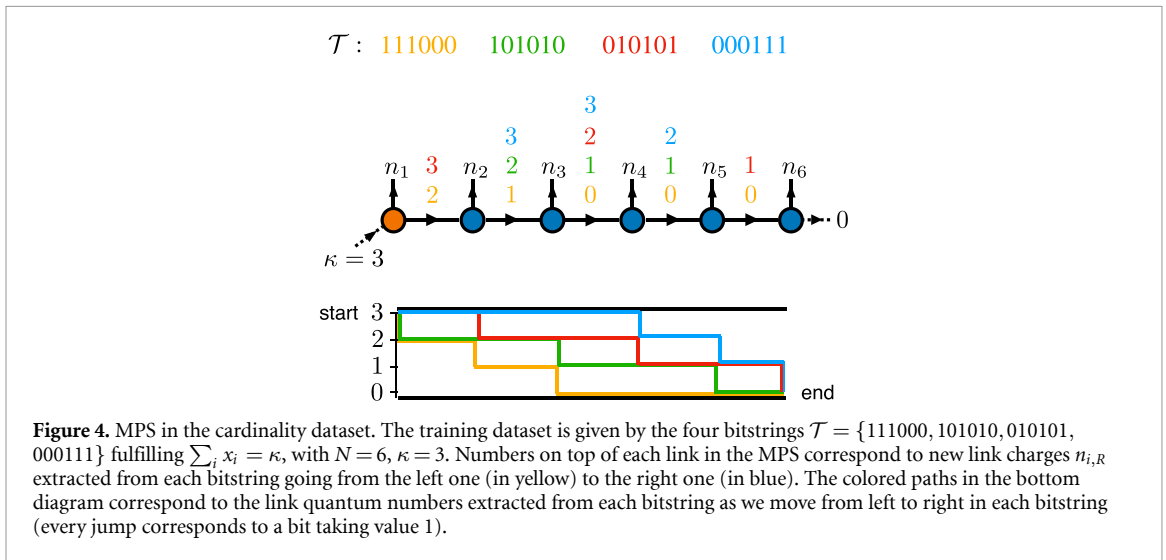
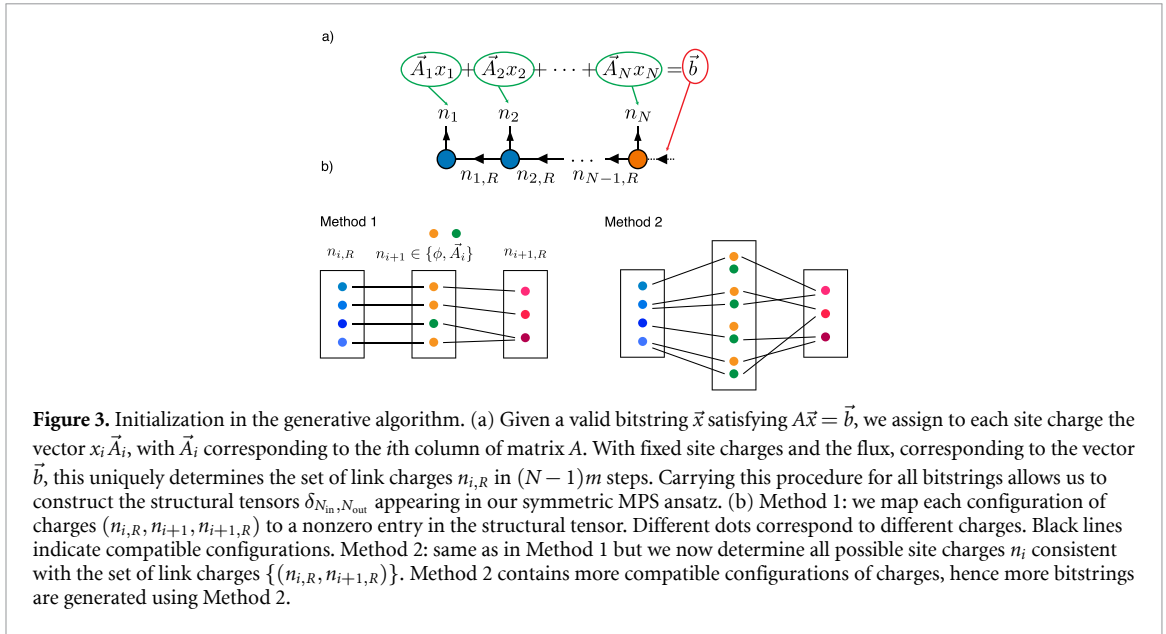
The most general case of arbitrary number of equalities of the form $A\vec{x} = \vec{b}$, with $A \in \mathbb{Z}^{m \times N}$, $\vec{b} \in \mathbb{Z}^m$, $\vec{x} \in \{0, 1\}^N$ is just a small extension of the previous case. Now, each site charge is instead given by a vector of length m as $n_i \in \{\emptyset, \vec{A}_i\}$ where $\vec{A}_i \equiv (A_{1,i}, A_{2,i}, \dots, A_{m,i})$, while the flux is given by the vector \vec{b} . In appendix C we consider ASSIGNMENT-type equalities which appear also in many combinatorial optimization problems and that can be encoded in an MPS of bond dimension scaling at worst exponentially with the number of equalities (i.e. these do not scale with system size N).

4. Constrained-GEO

Our framework is divided into three stages (see figure 1): (1) initialization, (2) training, and (3) sampling, followed by evaluation of the generalization performance. First, we construct the structural tensors based on the training data, or, if an efficient MPS representation of the valid space is available (such as in cardinality type constraints), we may also generate training data uniformly from it. Second, we minimize the training loss function via the DMRG-inspired algorithm put forward in [4]. After a few sweeps we sample from the resulting MPS and compute the costs of the samples, and retrain the model with those samples together with the initial training dataset, as we detail next.

4.1. Initialization: constrained embedding

The first step consists in selecting all unique bitstrings from the dataset. Next we construct the set of link charges by first fixing the position of the canonical center as well as that of the flux. For the purposes of illustration we set both to be at the last site as shown in figure 3(a). The determination of all link charges for each bitstring is performed in $\mathcal{O}(Nm)$ steps by solving recursively $n_{i+1,R} = x_{i+1}\vec{A}_{i+1} + n_{i,R}$ starting from either end of the MPS. We store in memory each local configuration of charges $i \rightarrow (n_{i,R}, n_{i+1}, n_{i+1,R})$, and move on to the next bitstring. We repeat the same procedure for the next bitstring in the set, and obtain a new configuration of charges. Once we loop over all distinct bitstrings from the training set we are ready to set up the structural tensors by reading off the configuration of charges. At this step, there are two alternative



ways of proceeding forward, as illustrated in figure 3(b). The first one is to directly map each configuration of charges $(n_{i,R}, n_{i+1}, n_{i+1,R})$ to a nonzero component of the structural tensor. Together with the finding of charges, this preprocessing step takes $\mathcal{O}(Nm|\mathcal{T}|)$ steps when using a hashtable for storing the configuration of charges. The second approach is to only store the configuration of all link charges consistent with the training data plus constraints, $i \rightarrow (n_{i,R}, n_{i+1,R})$, and subsequently determine all configurations of charges $(n_{i,R}, n_{i+1}, n_{i+1,R})$ consistent with the flux of each tensor. In contrast with the previous method, this preprocessing step scales as $\mathcal{O}(dNm|\mathcal{T}|)$ when using a hashtable, where d is the dimension of each site, which for our purposes will be 2 given that we are dealing with bitstrings. Thus we see that this latter method takes in general twice as much time as the first one. The benefit nevertheless from using this latter method is that this permits us to find more configuration of charges, which in turn will lead to better generalization, i.e. more (unique) bitstrings can be generated from the second method than with the first one.

To illustrate these two methods we consider the cardinality dataset formed by bitstrings fulfilling $\sum_i x_i = \kappa$. Focusing on bitstrings of length $N = 6$ and cardinality $\kappa = 3$ we choose the training data $\mathcal{T} = \{111000, 101010, 010101, 000111\}$. The resulting MPS using the first method will generate 6 new bitstrings (10 in total), while the second method will generate the entire solution space of $\binom{6}{3} = 20$ bitstrings. The extraction of link charges from data is illustrated in figure 4. There we show all possible QNs for each link index, $n_{i,R}$, consistent with the cardinality constraint. In total, there are $\binom{6}{3} = 20$ bitstrings. Out of them, only 4 are needed to extract all link charges. For arbitrary cardinality constraint κ and bitstring length N , it can be shown that only $\min(\kappa, N - \kappa) + 1$ bitstrings, are needed to exactly construct the fixed-cardinality MPS using the second method. In our work we shall use the second method when initializing the MPS.

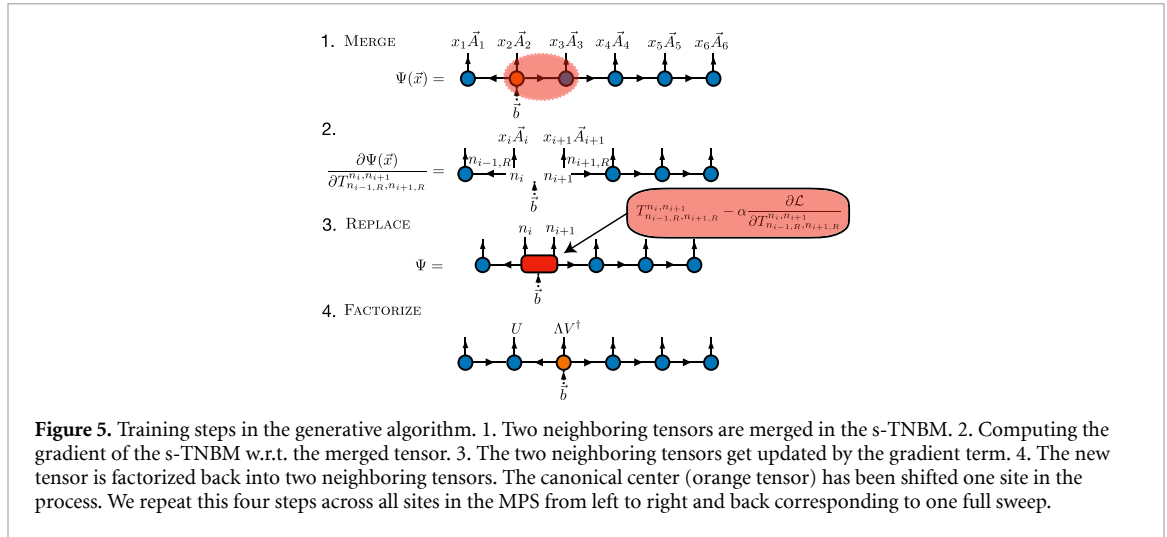


Figure 5. Training steps in the generative algorithm. 1. Two neighboring tensors are merged in the s-TNBM. 2. Computing the gradient of the s-TNBM w.r.t. the merged tensor. 3. The two neighboring tensors get updated by the gradient term. 4. The new tensor is factorized back into two neighboring tensors. The canonical center (orange tensor) has been shifted one site in the process. We repeat this four steps across all sites in the MPS from left to right and back corresponding to one full sweep.

We remark that the structural tensor crucially depends on the location where we insert the flux, which corresponds to the location of the canonical center. This is important to bear in mind because the training procedure will move the canonical center when optimizing each local tensor as we shall describe next (however the most costly step is setting up the structural tensors at the beginning, as updates during training in TNs are local, which will only change the link charges around each site, and these charge updates can be determined efficiently).

Lastly, we note that encoding arbitrary number of equality constraints (without training data) into a TN may lead to an exponential scaling of the bond dimension (see section C of the appendix for the particular case of ASSIGNMENT-type equalities). The present approach, guarantees that, in the worst case scenario, the bond dimension be lower bounded by the number of bitstrings in the training set, which allows for better scalability of the algorithm.

4.2. Training

Once we have set up the structural tensors deterministically, either via the initialization step described in section 4.1 or by exact construction of the MPS, we can proceed with the *training* step. The process is an extension of the DMRG inspired algorithm described in [4] in the presence of equality constraints. The training process aims to minimize the loss function

$$\mathcal{L} = - \sum_{\vec{x} \in \mathcal{T}} p_T(\vec{x}) \log \mathbb{P}(\vec{x}), \tag{6}$$

where $\mathbb{P}(\vec{x}) = |\Psi(\vec{x})|^2 / Z$, $Z = \sum_{\vec{x} \in \{0,1\}^N} |\Psi(\vec{x})|^2$, and p_T is a softmax surrogate to the cost function, and is given as in equation (2). Since the Born machine $|\Psi(\vec{x})|^2$ is constructed out of a symmetric MPS Ψ , we shall henceforth refer to this as symmetric-TNBM (s-TNBM). The various steps involved in the training are summarized in figure 5 and in more detail consist of the following:

1. For fixed location of the canonical center, construct the corresponding s-TNBM for a given bitstring $\vec{x} \in \mathcal{T}$ by fixing the site charges to $x_i \vec{A}_i$. Merge the two neighboring tensors at the canonical center. In TN notation:

$$\sum_{m_{i,R}} T_{n_{i-1,R}, m_{i,R}}^{[i] x_i \vec{A}_i} T_{m_{i,R}, n_{i+1,R}}^{[i+1] x_{i+1} \vec{A}_{i+1}} = \text{MERGE} \left(\begin{array}{c} x_i \vec{A}_i \quad x_{i+1} \vec{A}_{i+1} \\ \leftarrow \quad \rightarrow \\ \leftarrow \quad \rightarrow \\ \leftarrow \quad \rightarrow \\ n_{i-1,R} \quad n_{i,R} \quad n_{i+1,R} \\ \downarrow \\ \vec{b} \end{array} \right) = n_{i-1,R} \leftarrow \begin{array}{c} x_i \vec{A}_i \quad x_{i+1} \vec{A}_{i+1} \\ \leftarrow \quad \rightarrow \\ \leftarrow \quad \rightarrow \\ \leftarrow \quad \rightarrow \\ n_{i-1,R} \quad n_{i+1,R} \\ \downarrow \\ \vec{b} \end{array} \equiv T_{n_{i-1,R}, n_{i+1,R}}^{x_i \vec{A}_i, x_{i+1} \vec{A}_{i+1}}$$

where the link charges are constrained by charge conservation as $n_{i-1,R} + n_{i+1,R} = \vec{b} - x_i \vec{A}_i - x_{i+1} \vec{A}_{i+1}$. Because of charge conservation, the merged tensor will preserve block-sparsity.

2. Compute gradient of NLL as

$$\frac{\partial \mathcal{L}}{\partial T_{n_{i-1,R}, n_{i+1,R}}^{m_i, m_{i+1}}} = \frac{Z'}{Z} - 2 \sum_{\vec{x} \in \mathcal{T}} p_T(\vec{x}) \frac{\Psi'(\vec{x})}{\Psi(\vec{x})}, \tag{7}$$

In contrast, TNs enjoy the property of *perfect sampling* [25]. The bit-wise sampling of MPS is done without rejections and sequentially as one moves from one end of the MPS to the other, consisting on computing a series of single-site conditional density matrices and conditional single-site probabilities. The resulting sampling algorithm avoids long autocorrelation times typical of MCMC sampling-based algorithms that are needed to extract reliable samples. Since there is no correlation between samples, it is straight-forward to parallelize the sampling process by, for instance, making use of multi-threading.

Once samples are drawn, we evaluate the generalization performance of the algorithm by computing the costs of the bitstrings. The performance of the generative model will be measured according to (1) the *quantity* of new, valid samples (i.e. bitstrings not contained in the dataset, but that still fulfill the constraints), and (2) the *quality* of new samples as measured by the cost (higher-quality samples corresponding to lower-cost samples). These samples may be merged with the training dataset to form a new training set which we can be used for initialization and training in a new iteration of the algorithm. Various strategies can be devised for selecting this new training set as detailed in [21].

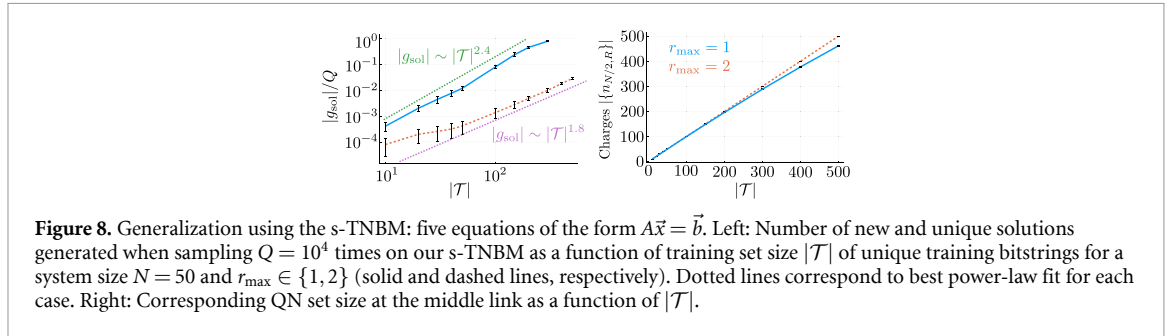
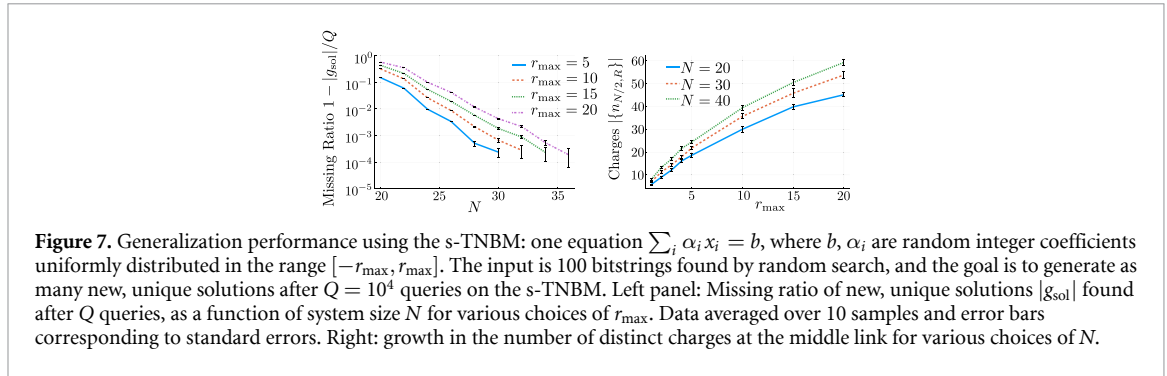
5. Results and discussion

To test the performance of s-TNBMs we consider the following tasks. In the first part, section 5.1, we will analyze how well symmetric MPS are able to generate new valid solutions to $A\vec{x} = \vec{b}$ from a given set of known solutions. This will test the ability of the structural tensors at capturing QNs that are shared among different bitstrings, which in turn will lead to generalization. In the second part, we will endow QN blocks with a degeneracy factor, taking full advantage of both QNs and their dimensions, so as to model arbitrary probability distributions (not just uniform on the space of solutions) subject to arbitrary equality constraints. We will test this on two tasks: first, in section 5.2, we shall consider the training dataset to be sampled from the ground state of a well-known Hamiltonian in condensed matter physics: the one dimensional transverse field Ising (TFI) model [39], subject to the constraint that only bitstrings of fixed cardinality are valid. Next, in section 5.3, we will test the ability of s-TNBMs at finding new, *high-quality* solutions to combinatorial optimization problems subject to equality constraints by mapping this problem to a generative modeling task, as briefly discussed in section 2, and detailed in section 4.

5.1. Validity-based generalization: entanglement assisted generation of solutions to arbitrary equality constraints

Evaluating the generalization performance of unsupervised generative models is a challenging task. For the following experiments, a candidate generalization metric should favor valid samples that are not only new, but unique as well. Our focus will be thus to measure variants of the set of new, and unique solutions, denoted as $|g_{\text{sol}}|$ (see [40] where this and other various metrics were introduced).

In the next experiments we shall be interested in generating new solutions to a system of linear constraints of the form $A\vec{x} = \vec{b}$, where $\vec{x} \in \{0, 1\}^N$, $A \in \mathbb{Z}^{m \times N}$, $\vec{b} \in \mathbb{Z}^m$ from a set of known solutions. In our first experiment we consider a single equation of the form $\sum_i \alpha_i x_i = b$, where α_i and b are random integers uniformly distributed in the range $[-r_{\text{max}}, r_{\text{max}}]$. We start with 100 unique known solutions to this equation, which we shall also call *seeds* (found e.g. by random search, or SUBSET SUM solvers when possible) and the goal is to generate as many solutions as possible. Since there is no bias in the training dataset (i.e. all solution bitstrings carry the same probability), we do not need to carry any explicit training, and instead we only make use of the initialization step explained in section 4.1, followed by the sampling step on the resulting s-TNBM. Since the solution space may grow with system size (this growth could be in principle exponential), we bound the number of solutions found by the number of total queries to the s-TNBM generator, which we set to be $Q = 10^4$. The results of this experiment are shown in figure 7. Remarkably we find that, despite the absence of obvious global $U(1)$ symmetry given the nature of the equality constraint, we do find very good generalization performance. It is interesting to see that the performance increases as we scale up system size N . The number of generated solutions grows at a pace which would be consistent with exponential (we do not rule out power law) as a function of system size N for fixed r_{max} , even for large values of r_{max} . Importantly, the set of QNs at the middle link, $|\{n_{N/2,R}\}|$ (which corresponds to the largest structural tensor link dimension in the s-TNBM), is of size at worst linear as a function of r_{max} for fixed N , at least for the system sizes and choices of r_{max} simulated here. Note this set size is a measure of the complexity of datasets constrained by equalities, and it provides a lower bound on the minimum bond dimension needed for an MPS to generate a given number of new, distinct solutions. Clearly, real-world equality constraints will have a greater degree of symmetry so the results shown here for a single equality with random coefficients A are already encouraging. This would in turn translate into a smaller bond dimension for a given fixed system size, and therefore better scalability as we go to larger system sizes.

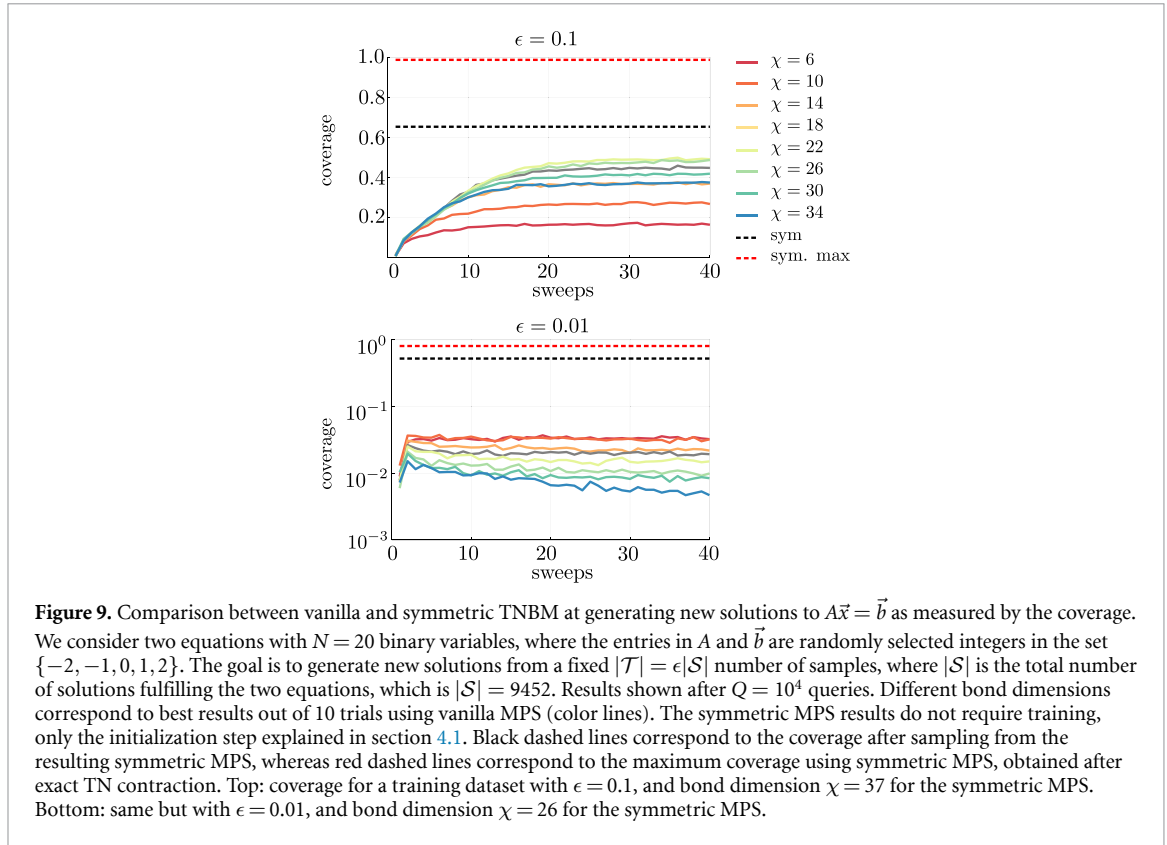


To have a better intuition on these results and on the role played by the bond dimension, it is helpful to consider the most *symmetric* equality constraint, that of cardinality, where $\alpha_i = 1 \forall i$. As argued in section 3.2, an exact MPS representation of the entire valid space requires a bond dimension of at least $\chi = \mathcal{O}(\kappa)$ (in fact $\chi = \min(\kappa, N - \kappa) + 1$), while the solution space grows exponentially in N (for $\kappa \approx N/2$). How does this relate to the generalization capabilities of s-TNBMs? In figure 4 we gave an intuition of this mechanism, focusing on $N = 6, \kappa = 3$. For concreteness we choose the canonical center to be located at the first site. We label all link indices by their QNs. Next, we apply the initialization step of the algorithm explained in section 4, with the set of four colored bitstrings in figure 4. These four samples already covers all possible QNs for each link with bond dimension 2, 3, 4, 3, 2 respectively. The corresponding MPS is able to generate all valid $\binom{6}{3} = 20$ bitstrings. The choice of the four samples is not unique as long as they *patch* the entire set of link QNs from start, at $n_{0,R} = \kappa$ (viewing the extra dashed index on the left of the MPS in figure 4 as having charge κ) to finish, at $n_{6,R} = 0$ (viewing the extra dashed leg on the right of the MPS in figure 4 as having charge 0). Of course this experiment carries over to any cardinality and system size. The upshot is that the minimum number of bitstrings needed to generate all solutions to a given equation is lower bounded by the number of QNs at the middle link. In other words, how well our s-TNBM is able to generalize depends crucially on the number of distinct paths that can traverse this middle link.

We scale up the previous experiment from one to five equations, where now A, \vec{b} entries are chosen uniformly random in $[-r_{\max}, r_{\max}]$. The results of the experiment are shown in figure 8. We are interested here in checking how efficiently does our s-TNBM generalize as a function of training set size, i.e. how many new, unique solutions, as a function of the training set size $|\mathcal{T}|$ are found by the s-TNBM. The training set is provided by random search and keeping valid samples. We find that, even for random instances as the ones considered in here, where $r_{\max} \in \{1, 2\}$, our s-TNBM is able to generate new solutions with a power-law exponent close to 2, as shown by the best-fit power-law exponent 1.8 and 2.4 in figure 8. The corresponding QN set size at the middle link grows linearly, and in fact scales roughly as $|\mathcal{T}|, |\{n_{N/2,R}\}| \approx |\mathcal{T}|$.

So far our focus has been solely on the generalization capabilities of s-TNBMs at generating new solutions to $A\vec{x} = \vec{b}$, without any comparison to the performance of vanilla TNBMs at the same task. In figure 9 we consider that comparison. A vanilla TNBM does not include any bias from the constraints, so one must train them using the DMRG-like algorithm from [4] (and extended to the presence of equality constraints in section 4.2). The experiment now considers a set of two equations with integer coefficients randomly selected in the range $[-2, 2]$. We are interested in the number of new, unique solutions, $|g_{\text{sol}}|$, found after a given number of queries Q to the generator. A useful metric in this sense is the coverage, as introduced in [40], and defined as

$$C = \frac{|g_{\text{sol}}|}{|\mathcal{S}| - |\mathcal{T}|}, \tag{8}$$

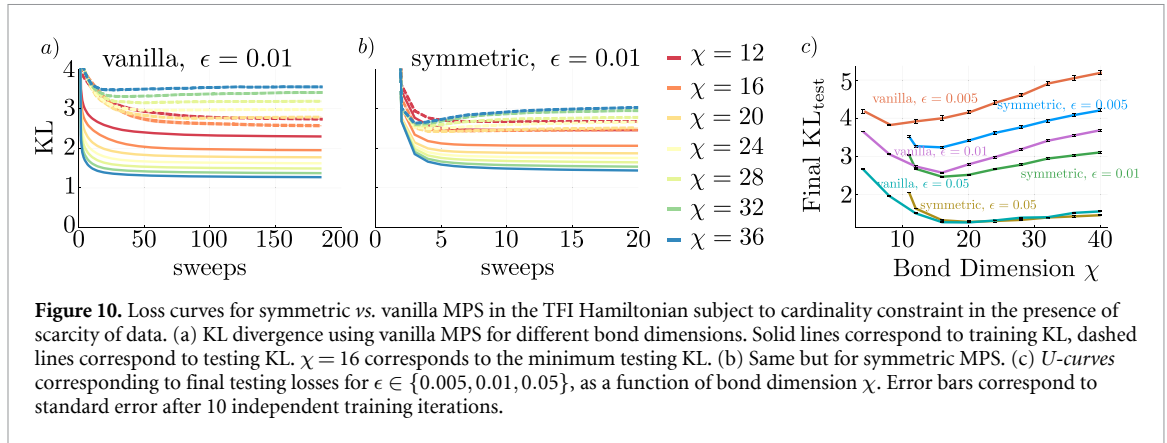


where $|\mathcal{S}|$ is the number of solutions to the system of equations. This metric measures effectively the number of new, unique solutions found by the generator within the remaining available solution space (not contained in the training set). Note that this quantity is implicitly a function of the number of queries to the generator. Given the system size chosen, $N = 20$, we can compute explicitly all solutions to random system of equations, which for our experiment turns out to be $|\mathcal{S}| = 9452$. We feed in both symmetric and vanilla MPS the same amount of training data, given as $\mathcal{T} = \epsilon|\mathcal{S}|$, with $\epsilon < 1$, which consist of unique, valid samples. When the training set contains 10% of the solutions ($\epsilon = 0.1$), both vanilla and symmetric MPS are able to generate new solutions with comparative performance. Two crucial differences between the vanilla and symmetric MPS results is that *a priori* we do not know which bond dimension gives the best performance using vanilla MPS and so one is left with trying many of them only to pick up the one that gives the best results, in this case $\chi = 22$. This is not the case when using symmetric MPS, where the bond dimension is fixed entirely by the training set and the constraints, i.e. it is given by the QN set size at the middle link, $\chi = |\{n_{N/2,R}\}|$. Perhaps more importantly, there is no training required when using symmetric MPS because the ansatz already contains the inductive bias coming from the constraints, and we only need the initialization step from section 4.1. The benefits of using symmetric MPS at learning datasets with equality constraints become more evident in the presence of scarcity of training samples: when we lower this amount down to 1% ($\epsilon = 0.01$), vanilla MPS produces about 3% of the remaining unique solutions after Q queries. This is true even for various choices of learning rates used in here, with the results of $\alpha = 0.02$ displayed in figure 9, whereas the symmetric version does so at around 50%, down from around 55% when $\epsilon = 0.1$.

5.2. Quality-based generalization: ground state reconstruction of local symmetric Hamiltonians

In this part we consider a generative model capturing equality constraints as well as a distribution bias. Our goal here is to learn the distribution corresponding to the ground state of the TFI model subject to an equality constraint which for simplicity we take it to be of cardinality type. The TFI Hamiltonian reads

$$\mathcal{H} = -s \sum_{i=1}^{N-1} Z_i Z_{i+1} - (1-s) \sum_{i=1}^N X_i, \quad (9)$$



where $s \in [0, 1]$, and $\{X_i, Z_i\}$ are Pauli matrices. Here we shall take $s = 1/2$, corresponding to the critical point of the TFI model, and the system size to be $N = 20$. This Hamiltonian by itself is not $U(1)$ symmetric (however it does have an inherent global \mathbb{Z}_2 symmetry, since the Hamiltonian is invariant under the parity operator $P = \prod_{i=1}^N X_i$), but its ground state is nevertheless efficiently captured by an MPS of low bond dimension, $\chi \approx 8$ (note however that for general N , the MPS ansatz is suboptimal at the critical point since entanglement grows logarithmically $S \sim \log N$, and instead the multi-scale entanglement renormalization ansatz is a better choice, see [41]). The target distribution corresponds to that constructed from taking the Born machine associated with this ground state followed by a *fixed cardinality filter*, that is, we project the resulting MPS onto the manifold of states constrained to only sample bitstrings fulfilling $\sum_i x_i = \kappa$, where for concreteness we choose the cardinality $\kappa = N/2 = 10$. Thus the entire solution space (composed of all valid bitstrings fulfilling the cardinality constraint) is $|\mathcal{S}| = \binom{20}{10} \approx 185 \times 10^3$.

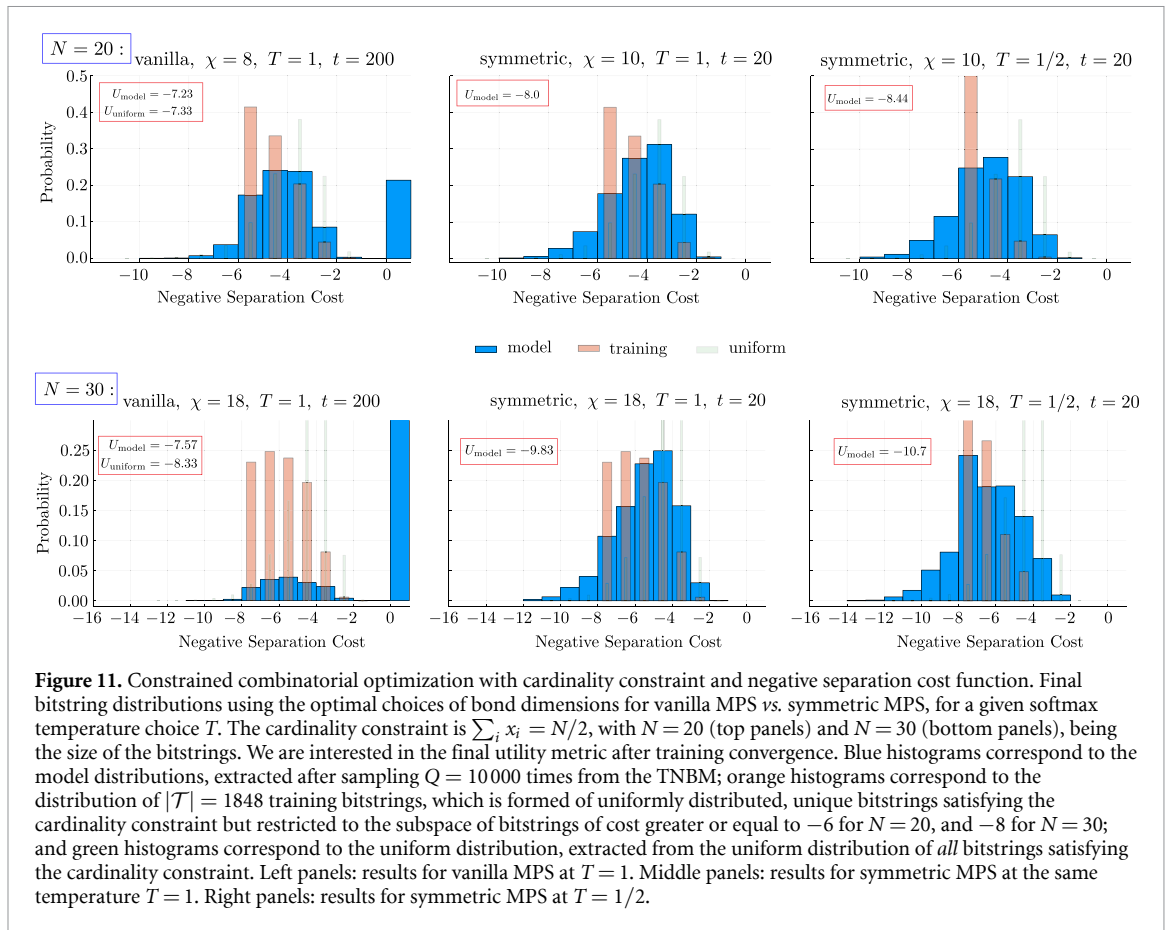
In figure 10 we show the training and testing loss curves in terms of their respective Kullback–Leibler divergences using vanilla vs. symmetric MPS for a training dataset size of $|\mathcal{T}| = \epsilon|\mathcal{S}|$ (note that the set may include repeated samples, in which case their probabilities are weighted according to their frequency). We see that while for $\epsilon = 0.05$ both vanilla and symmetric MPS are able to learn the underlying distribution just as well, when data is scarce, $\epsilon \lesssim 0.01$, the symmetric counterpart performs better. These results are obtained without major refinements, in particular the learning rate is fixed throughout and of value $\alpha = 0.02$ (in principle it would be advisable to adjust this as a function of the size of the tensors [4]), no DMRG noise is used, and no minibatching or multiple *inner steps* (i.e. multiple consecutive gradient updates at the same site) has been used. These tricks can bring down the value of the loss function by escaping local minima and may be crucial steps for larger scale models.

It is a well-known fact in machine learning (ML) models, that the testing loss may develop a *U-shaped* curve as a function of model capacity [42]. This is a consequence of the *bias-variance tradeoff*. In figure 10(right panel) we show that both vanilla and symmetric MPS display such a *U*-curve, as a function of bond dimension (see also [43]). This type of overfitting behavior is *static* in that it only depends on the complexity of the MPS as parameterized by its bond dimension χ . On the left and central panel in the same figure we also report a different kind of overfitting behavior which we dub *dynamic* in that it develops as a function of sweeps, and is more pronounced in the case of symmetric MPS, where at large enough χ , the best testing loss is localized at around a minimum after a few sweeps—such dynamic *U*-curve is also present in vanilla MPS, but its convexity is less striking.

We remark that we have also tried more suitable Hamiltonians with an inherent cardinality constraint, such as the Heisenberg model. This considerably improves the testing loss (since the Hamiltonian already satisfies the constraint) but we found the minimum in the *U*-curve to be located at the minimum bond dimension needed to learn the constraint, $\chi = 11$ for $N = 20$, which of course is not the general case for arbitrary cost functions (not just the ones stemming from local Hamiltonians).

5.3. Quality-based generalization: finding optimal solutions to constrained combinatorial optimization problems

The previous example was motivated by physics, in particular by exploiting the locality of the cost function in the form of the TFI Hamiltonian. Here we illustrate how s-TNBMs can also be used to find novel and optimal solutions to combinatorial optimization problems subject to equality constraints, outperforming



vanilla TNBMs. As an example, let us consider the cost function to be the *negative separation cost* of bitstrings, which is the negative separation of the farthest two ‘1’ separated only by ‘0’ bits. For instance, for a bitstring 01011001 the negative separation cost would be -3 . The goal of the task is to generate bitstrings with the largest negative separation cost (in absolute value) as possible, subject to an equality constraint which we take to be the cardinality constraint. This is achieved using the constrained-GEO framework described in section 4, where the cost function appearing in the softmax corresponds to the negative separation cost subject to a cardinality constraint. Because of this constraint, we assign a cost zero to invalid samples (bitstrings not fulfilling this constraint). The starting point is a set of valid solutions that have a negative separation cost at or above C_0 . We remark that as opposed to the previous task, with a physics motivation, the current exercise serves as a proxy for typical cardinality constrained optimization problems appearing in industry problems, such as portfolio optimization [12] and traveling salesman [44], both NP-hard problems [44, 45].

The results of this experiment are shown in figure 11, where the learning rate is fixed to $\alpha = 0.02$ throughout, and we choose two different system sizes. For $N = 20$ and $\kappa = 10$, the training set contains only bitstrings of cost greater than or equal to $C_0 = -6$, while for $N = 30$ and $\kappa = 15$, the training set contains only bitstrings of cost greater than or equal to $C_0 = -8$. To test the performance of both models at capturing low cost samples outside the training set we make use of the utility metric, as introduced in [40]. It amounts to computing the average cost of the 5% lowest cost samples obtained from querying the trained models. The motivation for this metric is that, often times, one is not interested in a unique lowest cost sample (which could have been found by mere luck). By taking instead a small percentage of the lowest cost samples, we make a more robust assessment that the model is indeed obtaining consistently high quality (low cost) samples. Our results indicate that, while the vanilla TNBM is not able to extract any better samples than a *constrained* uniform sampler, corresponding to the model that samples all bitstrings of the given cardinality uniformly, the symmetric version not only does sample already within the valid target space of fixed cardinality (thus avoiding samples with cost zero), but their samples are of higher quality as evidenced by the utility metric U . The value of this metric is even further optimized when decreasing the temperature in the




Figure 12. Flowchart of all the steps entering in the tensor network generative algorithm for combinatorial optimization problems with hard-constraints. See also algorithm 1 for a detailed description of the various steps.

softmax function, giving further evidence that the s-TNBM is able to learn the bias and that biasing the training set helps the s-TNBM to propose even better candidates. Crucially, we see that the performance gap between symmetric and vanilla TNBMs increase as we increase the system size from $N = 20$ to $N = 30$: in the latter case, the vanilla TNBM is not able to capture the cardinality constraint to the extent that less than 20% of the samples are valid, and more importantly, the quality of the valid ones, as measured by the utility, is still far from that obtained using s-TNBM.

Algorithm 1: Tensor network generative algorithm for hard constraints.

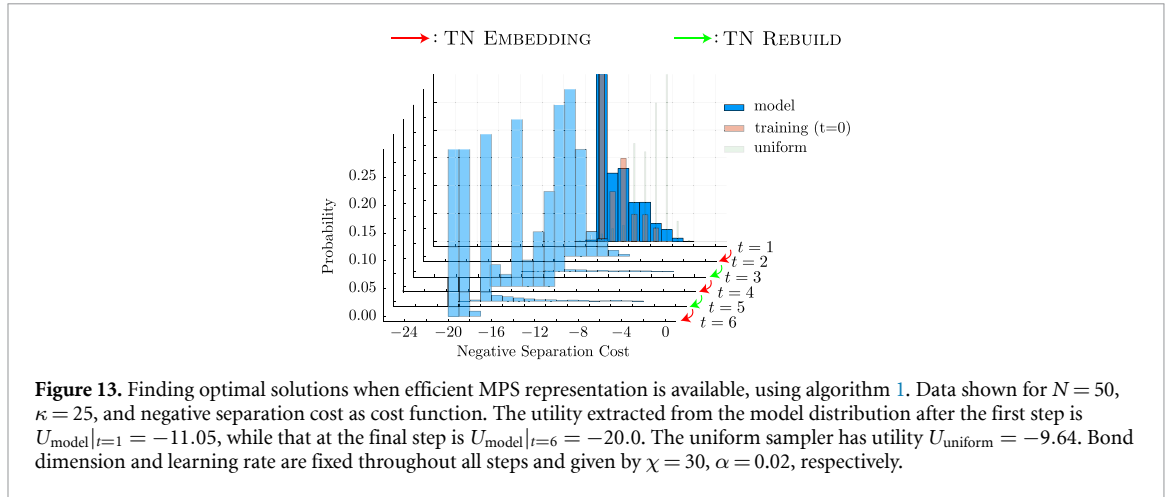
```

Input  Callback cost function  $\gamma = C(\vec{x})$ ,  $\vec{x} \in \{0, 1\}^{\otimes N}$ , and equality constraints  $A\vec{x} = \vec{b}$ 
Output  $\vec{x}^* \approx \operatorname{argmin}(C(\vec{x}))$ 
 $\psi(\vec{x}) \leftarrow \text{Constraints } A\vec{x} = \vec{b}$  ▷ Build* uniform TN from constraints (all QN blocks set to same value)
 $t \leftarrow 0$ 
 $\mathcal{T} \leftarrow \text{SAMPLE}(|\psi|^2)$  ▷ Extract training set via perfect sampling from s-TNBM
 $U|_t \leftarrow \text{mean}(5\% \text{ lowest } \{C(\vec{x}) : \vec{x} \in \mathcal{T}\})$  ▷ Compute utility from mean of lowest cost 5% samples
while ( $U|_{t+1} - U|_t > \epsilon$ ) do
   $t \leftarrow t + 1$ 
  if  $t$  odd then
     $\mathcal{T} \leftarrow \mathcal{T}|_{5\%}$  ▷ Create new training set from 5% lowest cost samples
     $p_T(\{\vec{x} \in \mathcal{T}\}) \leftarrow \{C(\{\vec{x} \in \mathcal{T}\}), T\}$  ▷ Construct training distribution  $p_T(\vec{x}) = e^{-C(\vec{x})/T} / \sum_{\vec{x} \in \mathcal{T}} e^{-C(\vec{x})/T}$ 
     $\psi(\vec{x}) \leftarrow \{\text{Constraints } A\vec{x} = \vec{b}, \text{Samples } \mathcal{T}\}$  ▷ Constrained TN embedding from constraints and samples
  else
     $p_T(\{\vec{x} \in \mathcal{T}\}) \leftarrow \{\vec{x} \in \mathcal{T}\}$  ▷ Construct uniform probability distribution over training set
     $\psi(\vec{x}) \leftarrow \text{Constraints } A\vec{x} = \vec{b}$  ▷ Rebuild* uniform TN from constraints (all QN blocks set to same value)
  end if
   $\psi \leftarrow \text{TRAIN}(\psi)$  ▷ Train MPS with distribution  $p_T$  for a few sweeps
   $\mathcal{T} \leftarrow \text{SAMPLE}(|\psi|^2)$  ▷ Extract training set via perfect sampling from s-TNBM
   $U|_{t+1} \leftarrow \text{mean}(5\% \text{ lowest } \{C(\vec{x}) : \forall \vec{x} \in \mathcal{T}\})$  ▷ Compute utility from mean of lowest cost 5% samples
end while
Return  $\vec{x} = \operatorname{argmin}(C(\vec{x}) : \vec{x} \in \mathcal{T})$ 

```

*: When an efficient TN representation of the solution space does not exist, an alternative construction is to search randomly samples satisfying the constraints and applying the Constrained TN embedding step on these. This will only construct an approximate TN to the entire solution space (note that in general the TN ansatz will contain more solutions than are fed in from a set of seeds, as argued in section 4.1).

In the next experiment we take full advantage of the knowledge that an exact and efficient representation of the valid space exists in terms of a symmetric MPS of fixed cardinality. We consider here a much larger system with $N = 50$ and cardinality $\kappa = 25$, which results in a solution space of size $|\mathcal{S}| = \binom{50}{25} \approx 1.26 \times 10^{14}$. We use the following protocol, which is also summarized in figure 12 for general equality constraints and detailed in algorithm 1. The first step is to sample uniformly from the fixed cardinality MPS. To guarantee that this sampling is uniform over the space of fixed cardinality bitstrings, we set all allowed QN blocks to be of size 1×1 and filled in by the same value (the value will be set by the normalization condition of the TN). This forms our initial training set. We train as usual for one sweep with this initial choice of MPS and sample at least as many times as samples there are in the training set. For our experiments the size of the training set is always fixed to be $|\mathcal{T}| = 100$, while the number of samples or queries on the generator is $Q = 10^4$ (we emphasize again that sampling on an MPS can be done very efficiently). With the samples at hand we keep the 100 lowest cost samples as the new training data set. At the next step we apply the initialization step



described in section 4.1 so as to extract the most relevant QNs from where we construct our next initial MPS. Next we train for one sweep with an effective temperature which we choose to be $T = 1/2\text{std}(C(\{\vec{x} \in \mathcal{T}\}))$, with std the standard deviation. We next sample Q times and keep the lowest cost 100 samples for our next step training dataset. We repeat this procedure for as long as the results improve. The results of this experiment are shown in figure 13. After only 6 steps (and only a couple minutes of computational time on a regular laptop), the symmetric MPS is able to reach a minimum cost of $\mathcal{C} = -20$, with the absolute minimum corresponding to $\mathcal{C} = -26$. To put this result into context, we remark that the space of bitstrings with cost $\mathcal{C} \leq 20$, satisfies $|\mathcal{S}|_{\mathcal{C} \leq 20}/|\mathcal{S}| \approx 10^{-7}$ (see appendix D for an exact analytical expression), so that a constrained uniform sampler over the solution space would require at least of the order of $\mathcal{O}(10^3)$ times more samples than the ones carried out here, indicating that our MPS indeed is able to find its way to a subspace near the global minimum, even when this effectively corresponds to a tiny fraction of the entire valid space, and in turn the cost function is highly nonlocal.

This optimization scheme bears some resemblance with the QAOA and other approaches following an *exploration/exploitation* paradigm: the TN REBUILD step serves as a mixer near the optimum solution (which learns over the best data from the previous iteration), while the TN Embedding step, moves directly along the direction of minimum by learning only the most relevant QNs (*latent features*) from the lowest cost samples from the previous iteration. Our approach, being data driven, is agnostic to the form of the cost function, as illustrated in our experiments.

6. Comparison to neural network architectures

TNs are an efficient representation of states with limited amount of entanglement. In recent years there has been a lot of excitement about the possibility of using other representations that may be more efficient than TNs at capturing long-range entanglement. RBMs, the building block of many deep neural network architectures, have been suggested to be a more efficient representation than TNs for certain classes of states displaying long-range entanglement [46, 47]. Here we argue that $U(1)$ -symmetric states may be more challenging to construct using RBMs, and likely the same true for any \mathcal{G} -symmetric state, with \mathcal{G} any global internal symmetry, when compared to TNs. The intuition behind this argument relies on the fact that TNs are multi-linear models. In particular, the contraction of an MPS with one-site operators can be performed locally. Under a global group transformation represented by a unitary matrix $U(g)^{\otimes N}$, with $g \in \mathcal{G}$, the new MPS can be constructed straight-forwardly, see figures 14(a) and (b). In contrast, for a RBM, after applying such global group transformation, it is very challenging to get the weights of the new RBM. Unlike for a TN, we have to train the new RBM. So it becomes natural to impose such symmetries in TNs, in contrast to RBMs or ANNs in general.

For a TN, group operations on sites will introduce group representations on each link consistent with charge conservation. So a global symmetry transformation will generate group representations on each link of the MPS. While for RBM or ANN, it is very challenging to have this property.

Another argument in favor of TNs for describing symmetric states is precisely due to the linearity of the TN. As shown in figures 15(a) and (b), it is very easy to construct the linear hybridization of two MPS, the new tensors are a direct sum of these MPS tensors. The bond dimension would be the sum of bond dimensions of $\Psi_1(x)$ and $\Psi_2(x)$, and the construction is exact. In contrast, for a RBM, we are not able to

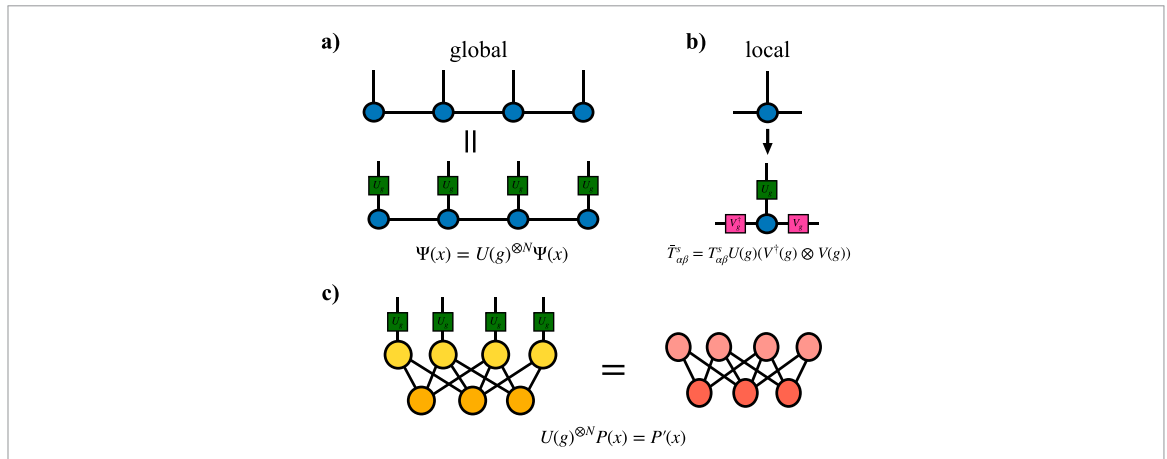


Figure 14. Illustration of how a TN and RBM are transformed under global symmetry. (a) $\Psi(x)$ is an MPS with \mathcal{G} symmetry. The MPS remains constant under a global rotation for all sites. (b) The local tensors can be associated with the original ones by contracting local tensor T with gate U . (c) Global symmetry transformation on a RBM, which results in a new RBM with updated weights. Due to the sparsity structure of the RBM network, it becomes very challenging to get the updated RBM parameters after the symmetry transformation.

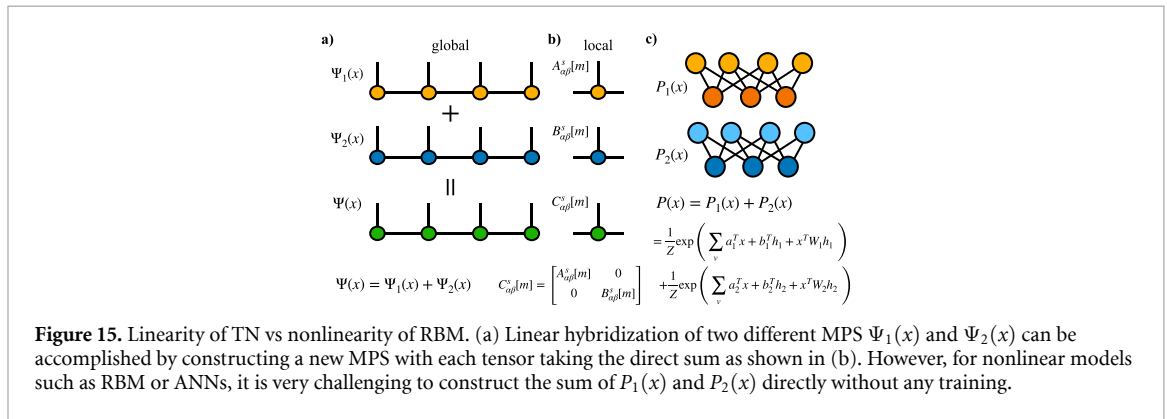


Figure 15. Linearity of TN vs nonlinearity of RBM. (a) Linear hybridization of two different MPS $\Psi_1(x)$ and $\Psi_2(x)$ can be accomplished by constructing a new MPS with each tensor taking the direct sum as shown in (b). However, for nonlinear models such as RBM or ANNs, it is very challenging to construct the sum of $P_1(x)$ and $P_2(x)$ directly without any training.

construct the new RBM directly from the original RBM parameters. The only option left is to train a new RBM, which mimics the summation, which is only an approximation.

7. Conclusions

In this work we have shown how one can use symmetric TNs for the purpose of generative modeling and constrained combinatorial optimization. We have tested their performance on various datasets serving as proxies for the kinds of problems our generative algorithm, dubbed here constrained-GEO, is applicable to. Most prominently, we include combinatorial optimization problems subject to arbitrary integer-valued equality constraints, which appear in many problems of industrial interest. We have found that by mapping the problem of equalities to a problem involving $U(1)$ global charges, a symmetric MPS generative model is able to generalize (even in the presence of random instances of A and \vec{b}). This class of models, referred here as s-TNBMs, has a series of benefits which we recount:

- By exploiting $U(1)$ global charge conservation, we are able to reduce the space complexity by working with block-sparse tensors, as opposed to dense tensors, as shown on the right panel of figure 6.
- In turn, given the block-sparse structure of tensors allows for faster learning of the data, often requiring an order of magnitude less number of sweeps to converge when training for the problem sizes studied in here.
 - * This point translates into less computational time per operation (contraction, merging) as shown on the left panel of figure 6;
 - * Less number of sweeps as shown on figures 10(a) and (b); and
 - * No time is spent on sampling invalid samples.

- It permits to have a better intuition on the complexity of the data, as viewed when counting the number of QN blocks appearing in the MPS (i.e. the minimum bond dimension needed to exactly capture all charges from the data).
- It allows to generalize in the presence of scarcity of data.
- It allows to find better solutions in the context of combinatorial optimization problems subject to constraints.

Of all these points, the last two ones are perhaps the most relevant ones. Classical ML models require many samples for them to generalize. By exploiting internal global symmetries like $U(1)$, which in turn are hard to impose in ANNs architectures as argued in section 6 (see also [48, 49]), we are able to reduce the amount of training data even further. The last point is also worth stressing: symmetric TNs were originally conceived to efficiently encode quantum states subject to global symmetries, but in the context of finding ground-states of Hamiltonians, as well as in the context of Hamiltonian dynamics. What we have found in this work is that constraining TNs may be a key step at not only decreasing the complexity of the learning algorithm, but crucially, at finding novel and better solutions to constrained combinatorial optimization problems, which would be otherwise out-of-reach by other classical ML and state-of-the-art optimization heuristics.

8. Outlook

A clear future extension of the current work is to analyze the role of other symmetries. Trivial extensions include other Abelian symmetries, like \mathbb{Z}_2 , which appear in physical models such as the TFI model discussed in this work, but also in some toy datasets like the *parity* dataset, for which TNs have already shown superior performance over ANNs such as RBMs at both learning and generative modeling performance [50–54]. Perhaps more intriguing would be to study the role of non-Abelian symmetries in the context of generative modeling. Of special relevance is $SU(2)$, perhaps the other most relevant *quantum*-inspired symmetry apart from $U(1)$, and for which a way to construct TNs that are symmetric under $SU(2)$ already exists in the literature [55–57]. A more challenging and practically relevant question is to address the presence of inequality constraints (i.e. of the type $A\vec{x} \leq \vec{b}$), which also abound in many combinatorial optimization problems. This is an ongoing work and will be presented in a subsequent publication.

We now put our work into the broader context of current progress in the domains of constrained combinatorial optimization and ML models exploring the role of symmetries and constraints. In the domain of combinatorial optimization, GNNs have become a very promising framework (see e.g. [17, 58]). In this case, one can encode the coefficients of matrix A and vector \vec{b} as features in a bipartite graph. This approach leverages generative modeling for *node selection* and *variable selection*, steps that appear in branch-and-bound for solving MILP problems. In the context of quantum models addressing QUBO problems with hard-constraints, such as XY-QAOA [59, 60], current methods suffer from the fact that only a very limited set of equality constraints can be addressed, namely those for which one can implement *local* moves around the valid space (and which can be implemented by at most two-site, local gates), such as cardinality-type constraints. Likewise, a recent TN approach was proposed to tackle the open-pit problem subject to equality constraints [8]. While bearing some resemblance to our work, this approach however is only valid for equality constraints with a locality structure, in that only few local variables can appear at any time in any equality, and where in turn the cost function must be mapped to a local Hamiltonian. Our approach is not limited by specific equality constraints nor by specific choices of cost functions.

On the side of implementing symmetries in the context of ML, there has been very recent interest in equivariant ANNs, both classical (see e.g. [61, 62]), and quantum (see e.g. [63–66]). Equivariant ANNs exploit symmetry in data so that the output samples also preserve this symmetry, such as is the case of $E(3)$ equivariant ANNs (which transform data under translations and rotations, with applications including image processing). In contrast, our work exploits symmetry at the level of the generator, in this case a s -TNBM, so that it only outputs samples satisfying arbitrary equality constraints, where different equality constraints correspond to different irreps of $U(1)$ (i.e. different fluxes in TN terminology). Lastly, there has been recent interest in using autoregressive neural networks for modeling physical systems with local gauge symmetries [67, 68], which are fundamentally different from the global internal symmetries we are interested in here. The latter have been studied in works [69, 70] for the simplest $U(1)$ symmetry corresponding to an equality of cardinality type, using recurrent neural networks (RNNs). However at present it is not clear how to extend their construction to arbitrary equality constraints, and if so, whether the resulting RNN corresponds to a more efficient representation than that of symmetric MPS.

An exciting milestone ahead is to demonstrate the power of our s -TNBM models in the context of large-scale industrial applications. Achieving such quantum-inspired advantage would be a promising step towards demonstrating practical quantum advantage with near-term devices. This could be possible by

exploiting recent efficient proposals [71, 72] to map TNs to quantum circuit generative models such as quantum-circuit Born machines [73].

Data availability statement

The data cannot be made publicly available upon publication because no suitable repository exists for hosting data in this field of study. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

The authors would like to acknowledge Mohamed Hibat-Allah, Marta Mauri, and Artem Strashko for helpful discussions and insights. We also acknowledge Dax Enshan Koh and Brian Dellabetta for reviewing our manuscript and providing feedback.

Appendix A. Introduction to (0,1)-INTEGER LINEAR PROGRAMMING

(0,1)-INTEGER LINEAR PROGRAMMING problems are optimization problems that involve binary decision variables, taking values 0 or 1. The objective is to find the combination of variables that maximize or minimize a given linear function subject to a set of linear constraints.

A.1. Formulation of the problem

The general form of a (0,1)-INTEGER LINEAR PROGRAMMING problem can be stated as follows [22]:

$$\begin{aligned} & \text{MAXIMIZE OR MINIMIZE } \vec{c} \cdot \vec{x} \\ & \text{SUBJECT TO } A\vec{x} \leq \vec{b}, \end{aligned}$$

where c_i and A_{ij} are integer coefficients of the objective function and the constraints, respectively. The variables x_i , take on binary values of 0 or 1, and b_i are integers. This class of problems belong to one of Karp's list of 21 NP-COMPLETE problems. Hence, solving instances of this problem takes in general exponential computational resources in the problem size.

There exist extensions of the (0,1)-INTEGER LINEAR PROGRAMMING to the case when some of the variables are let to be arbitrary reals and the rest arbitrary integers, in which case this class of problems goes by the name of MIXED INTEGER LINEAR PROGRAMMING (MILP) problems. This class of problems are well-known to be NP-HARD as well in general. Other extensions include the case when the objective function and/or constraints are not linear functions (and some variables are still integer valued). The complexity class corresponding to *nonlinear* integer programming is also NP-HARD and expected to be in general much harder to solve than their linear counterparts.

We will briefly discuss some standard methods for solving (0,1)-INTEGER LINEAR PROGRAMMING problems (we refer the reader to reference [74] discussing some of their extensions to the presence of nonlinear cost functions and linear constraints, which is the class of problems our TN algorithm is able to handle).

A.2. Solution methods

The solution to a (0,1)-INTEGER LINEAR PROGRAMMING problem can be found using a variety of methods. We distinguish two classes of methods: *exact* methods, including branch and bound, cutting plane, and branch and cut algorithms; as well as *heuristic* methods, including simulated annealing, parallel tempering, and genetic algorithms. The difference between these two classes is that the former provides guarantees of optimality of the solutions found, but do not scale well with problem size in general. The second class of methods, being heuristic do not have guarantees of optimality but do not suffer from the same scalability issues. Here we review the former class of methods since they are the ones most frequently used in most MILP solvers [14]. This class of methods involve iteratively solving a series of linear programming problems, where the decision variables are not restricted to binary values. All of these methods are discussed in detail in standard references such as reference [22].

Branch and bound –. The branch and bound algorithm works by dividing the search space into smaller subspaces (or *branches*), and using bounds on the objective function to eliminate subspaces that are guaranteed to contain suboptimal solutions. The algorithm then selects the subspace with the best solution and divides it again until the optimal solution is found.

Cutting plane –. The cutting plane algorithm involves adding additional linear constraints to the problem to exclude certain solutions. These additional constraints are based on the solution to the current linear programming problem and are designed to tighten the feasible region of the search space.

Branch and cut –. The branch and cut algorithm combines the techniques of branch and bound and cutting plane. The algorithm begins by using branch and bound to divide the search space, and then uses cutting plane to tighten the feasible region of each subspace.

Appendix B. Symmetric TNs

In this appendix we give a brief introduction to MPSs and their $U(1)$ symmetric versions, as done in [33, 34].

B.1. Vanilla TNs and MPSs

Consider a system of N sites where each local site is characterized by a local complex vector space \mathbb{V} . A wavefunction for this system is given as

$$|\Psi\rangle = \sum_{i_1, i_2, \dots, i_N} (\Psi)_{i_1, i_2, \dots, i_N} |i_1, i_2, \dots, i_N\rangle, \quad (\text{B1})$$

with $|i_v\rangle \in \mathbb{V}$. The coefficient $(\Psi)_{i_1, i_2, \dots, i_N}$ is in fact a *tensor*, of rank N . The main idea behind TNs is to decompose this big tensor into tensors of smaller rank contracted with each other such that [34]

$$(\Psi)_{i_1, i_2, \dots, i_N} = \text{t Tr} \left(\bigotimes_{v=1}^N T^{[v]i_v} \right), \quad (\text{B2})$$

where t Tr stands for tensor trace which contracts all repeated tensor indices of tensors $T^{[v]}$. A TN of special relevance is the MPS which is given as

$$(\Psi)_{i_1, i_2, \dots, i_N} = \text{Tr}(T^{[1]i_1} T^{[2]i_2} \dots T^{[N]i_N}), \quad (\text{B3})$$

where the trace now is a conventional matrix trace, where each tensor $T^{[v]i_v}$ is a $\chi_{v-1} \times \chi_v$ matrix with components $T_{l_{v-1}, l_v}^{[v]i_v}$. For our purposes, we shall be considering the MPS with *open boundary conditions* [75] so that the left- and right-most tensors are actually vectors so that $\chi_0 = \chi_N = 1$ (so that the trace is not needed). A perspective that is useful when dealing with symmetry group transformations on tensors, is to view tensors as linear maps from an *input* vector space \mathbb{V}^{in} to *output* vector space \mathbb{V}^{out} . A significant simplification in many TN algorithms involving MPS is to use the canonical form [24], which is used throughout in this work. When put in canonical form, tensors to the left of the canonical center correspond to left isometries, satisfying $|\beta\rangle = \sum_{a, \alpha} T_{\alpha, \beta}^a |a, \alpha\rangle$ with $\sum_a (T^a)^\dagger T^a = \mathbb{1}$ and where $|a, \alpha\rangle \in \mathbb{V}^{\text{out}}$ and $|\beta\rangle \in \mathbb{V}^{\text{in}}$, while tensors to the right of the canonical center correspond to right isometries, satisfying $|\alpha\rangle = \sum_{a, \beta} T_{\alpha, \beta}^a |a, \beta\rangle$ with $\sum_a T^a (T^a)^\dagger = \mathbb{1}$ and where $|\alpha\rangle \in \mathbb{V}^{\text{in}}$ and $|a, \beta\rangle \in \mathbb{V}^{\text{out}}$. Each *link* or *virtual* state can be represented as $|l_v\rangle \in \{|1\rangle, |2\rangle, \dots, |\chi_v\rangle\}$. The size of each link vector space is sometimes known as bond dimension, and its magnitude controls the amount of entanglement shared by connected tensors. The states associated to the upper index i_v will be referred to as *physical* states, since they are the ones that appear in the original wavefunction expansion (B1), at variance with the virtual states just discussed.

B.2. $U(1)$ symmetric MPSs

Since tensors are linear transformations, the TN ansatz makes it very convenient to exploit the toolbox of representation theory. In essence, by requiring that our state $|\Psi\rangle$ be *invariant* under a *global* group transformation \mathcal{G} represented by a unitary matrix U_g , with $g \in \mathcal{G}$ s.t. [33]

$$(U_g)^{\otimes N} |\Psi\rangle = |\Psi\rangle, \quad (\text{B4})$$

we may constrain the form of each tensor in the MPS as dictated by the transformation

$$\sum_{a', b', c'} (U_g)_{c, c'} (V_g)_{b, b'} T_{a', c'}^{b'} (W_g)_{a', a}^\dagger = T_{a, c}^b, \quad (\text{B5})$$

where V_g, W_g are unitary representations of the same group \mathcal{G} . The upshot of this is that each tensor will be constrained to have nonzero entries only at those entries fulfilling equation (B5). Our focus here will be on tensors that are *symmetric* under $\mathcal{G} = U(1)$, where representations are one dimensional and parameterized as $U_g = e^{-i\hat{n}\phi}$, with \hat{n} the so-called *charge* operator and $\phi \in [0, 2\pi)$. This means that our state must be invariant under this group transformation up to a phase:

$$e^{-i\hat{n}\phi} |\Psi\rangle = e^{-in\phi} |\Psi\rangle, \quad (\text{B6})$$

where the state $|\Psi\rangle$ is an eigenstate of the charge operator \hat{n} with eigenvalue n ,

$$\hat{n}|\Psi\rangle = n|\Psi\rangle. \tag{B7}$$

Note that, the case $n = 0$ corresponds to the state being invariant (as in (B4)), while for $n \neq 0$ one obtains a *covariant* state. Such phase factor only arises for states that are covariant under Abelian group transformations like $U(1)$ (in the case of non-Abelian group transformations, covariant states transform instead under a unitary matrix of dimension 2 or higher). The different charges n label different irreps of $U(1)$ so that for each vector space \mathbb{V} we may decompose this as

$$\mathbb{V} \simeq \bigoplus_n \mathbb{V}_n, \tag{B8}$$

where \mathbb{V}_n are irreps of fixed charge n , each of dimension (or *degeneracy*) d_n . Thus, any arbitrary vector $|\nu\rangle \in \mathbb{V}$ can be expanded as a linear combination of $|n, t_n\rangle \in \mathbb{V}_n$ vectors of fixed charge and degeneracy label t_n , where $t_n = 1, \dots, d_n$, with d_n the degeneracy of charge n . By going to this basis of well-defined charge, the invariance condition on the tensors, (B5), translates into

$$T_{a,c}^b = (T_{n_a, n_c}^{n_b})_{t_{n_a}, t_{n_c}}^{t_{n_b}} \delta_{N_{in}, N_{out}}, \tag{B9}$$

where

$$N_{in} = \sum_{i \in \mathcal{I}} n_i, \quad N_{out} = \sum_{i \in \mathcal{O}} n_i, \tag{B10}$$

where \mathcal{I} (\mathcal{O}) denotes the set of incoming (outgoing) indices. From (B9) we see that each tensor in the MPS factorizes into the product of two tensors. The first one, $(T_{n_a, n_c}^{n_b})_{t_{n_a}, t_{n_c}}^{t_{n_b}}$ gives the components of tensor T w.r.t. the degeneracy vectors $\{|t_n\rangle\}$, reason why this tensor goes by the name of *degeneracy tensor*. The second tensor, $\delta_{N_{in}, N_{out}}$ depends solely on the conservation of $U(1)$ charge, and it is referred to as *structural tensor* (for an arbitrary symmetry group \mathcal{G} the structural tensor can take a rather nontrivial form). For fixed physical charge (fixed n_b appearing in (B9)), the resulting matrix $T_{a,c}^b$ will be block diagonal, with blocks of size $d_{n_a} \times d_{n_c}$. These blocks are sometimes referred to as QN blocks. A symmetric MPS will be in general not only composed of invariant tensors fulfilling (B9) but also covariant tensors of well defined charge $n \neq 0$ and fulfilling

$$\sum_{a', b', c'} (U_g)_{c, c'} (V_g)_{b, b'} T_{a', c'}^{b'} (W_g)_{a', a}^\dagger = e^{-in\phi} T_{a, c}^b. \tag{B11}$$

Covariant tensors can be most compactly written in the basis of well-defined charge as

$$T_{a, c}^b = (T_{n_a, n_c}^{n_b})_{t_{n_a}, t_{n_c}}^{t_{n_b}} \delta_{N_{in} + n, N_{out}}. \tag{B12}$$

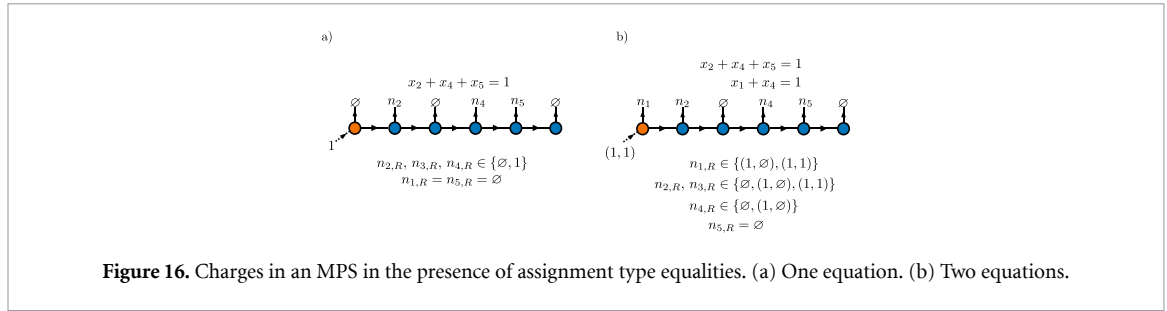
Alternatively, we may transform a covariant tensor to an invariant tensor upon introducing an extra index carrying charge n and of dimension 1. The charge n in a covariant tensor is also sometimes known as the *flux*. For a given MPS, there is at least one covariant tensor, and this is usually taken to be located at the canonical center of the MPS. Alternatively, we may decide to spread the total charge across different sites. It is however preferred to concentrate all flux in one single site as it is easier for bookkeeping purposes (we do not have to keep track of which sites have flux and which do not) and the resulting complexity (as measured by the size of the tensors) is not affected by this choice.

Lastly, we remark that for the purposes of our work where we have multiple arbitrary equality constraints, the above formalism still carries through. The major difference is that now we require the state to be invariant under multiple global group transformations of the form $U_g = e^{-i \sum_j \alpha_j \hat{n}_j \phi}$. Since each of the operators \hat{n}_j acts locally on each tensor, it is straightforward to generalize the above formulas to the case of arbitrary number of charges (i.e. translation invariance is not a requirement).

Appendix C. Exact MPS construction for assignment type problems

In section 3.2 we gave an exact construction of the MPS in the presence of a cardinality constraint. Here we give the derivation for another widely present equality constraint in combinatorial optimization problems, namely, an *assignment* type constraint. This is of the form

$$\sum_{i \in \mathcal{A}} x_i = 1, \tag{C1}$$



with \mathcal{A} some set of site indices. The valid space for this constraint can be encoded in an MPS of $\chi = 2$, with link charges $n_{i,R} \in \{\emptyset, 1\}$. See figure 16 for an illustration for a single equality. For m number of assignment type equalities, the link charges take values in $n_{i,R} \in \{\emptyset, 1\}^m$, see figure 16 for an example of two equalities. Thus at worst the bond dimension grows exponentially with the number of assignments.

Appendix D. Degeneracy counting for negative separation cost function

The negative separation cost function together with the cardinality constraint as introduced in section 5.3 has $\kappa - 1$ bitstrings with minimum cost function, which is $\mathcal{C}_{\min} = -N + \kappa - 1$. This corresponds to having a single connected block of zeros of size $N - \kappa$, with ones at each edge. For the remainder we shall refer to an arbitrary-sized block of zeros with ones at each side as a *domain wall* (to connect with the statistical physics literature). Here we wish to characterize the degeneracy (that is, number of bitstrings) for arbitrary value of the cost function and for any system size. This will be useful as a performance benchmark of the trained s-TNBM against our baseline which is a uniform sampler on the space of bitstrings with the right cardinality κ . Since for large system sizes $N \gtrsim 30$ it becomes numerically challenging to extract the degeneracy factors, we aim at extracting these analytically.

The problem of counting such configurations is a straightforward problem of combinatorics. The approach we take here is recursive. We present results for the case when there is a single dominant domain wall, meaning one domain wall is bigger than the rest in a given bitstring (this is analogous to a low-energy approximation above the absolute minimum when viewing the cost function as an energy function). Our results are as follows.

Let $a \in \mathbb{Z}$ with $2a < N - \kappa$, where κ is the Hamming weight (cardinality). The degeneracy in the number of bitstrings with cost $\mathcal{C} = -N + \kappa + a - 1$ is given as

$$|\mathcal{S}|(a, \kappa) = \begin{cases} 2 \sum_{i=a}^{\kappa+a-2} \binom{i}{a} + 2 \sum_{j=1}^{\frac{a-1}{2}} \sum_{i=a-j}^{\kappa+a-2-j} \binom{i}{a-j} \binom{\kappa+a-2-i}{j}, & \text{if } a \in 2\mathbb{Z} + 1, \\ 2 \sum_{i=a}^{\kappa+a-2} \binom{i}{a} + 2 \sum_{j=1}^{\frac{a}{2}-1} \sum_{i=a-j}^{\kappa+a-2-j} \binom{i}{a-j} \binom{\kappa+a-2-i}{j} + \sum_{i=\frac{a}{2}}^{\kappa-2+\frac{a}{2}} \binom{i}{\frac{a}{2}} \binom{\kappa+a-2-i}{\frac{a}{2}}, & \text{if } a \in 2\mathbb{Z}. \end{cases} \quad (\text{D1})$$

The way to derive this result consists in a simple counting of the number of configurations of zeros at each side of the main domain wall. In particular, for the experiment of figure 13, we have $|\mathcal{S}|(6, 25) / |\mathcal{S}| = \mathcal{O}(10^{-7})$, where $|\mathcal{S}| = \binom{50}{25}$ the number of solutions of fixed cardinality.

ORCID iDs

Javier Lopez-Piqueres <https://orcid.org/0000-0002-5188-9589>

Jing Chen <https://orcid.org/0000-0003-0538-689X>

Alejandro Perdomo-Ortiz <https://orcid.org/0000-0001-7176-4719>

References

- [1] Orús R 2019 Tensor networks for complex quantum systems *Nat. Rev. Phys.* **1** 538–50
- [2] Chen D-T, Helms P, Hale A R, Lee M, Li C, Gray J, Christou G, Zapf V S, Chan G K-L and Cheng H-P 2022 Using hyperoptimized tensor networks and first-principles electronic structure to simulate the experimental properties of the giant {Mn84} torus *J. Phys. Chem. Lett.* **13** 2365–70
- [3] Stoudenmire E and Schwab D J 2016 Supervised learning with tensor networks *Advances in Neural Information Processing Systems* vol 29
- [4] Han Z-Y, Wang J, Fan H, Wang L and Zhang P 2018 Unsupervised generative modeling using matrix product states *Phys. Rev. X* **8** 031012
- [5] Wang J, Roberts C, Vidal G and Leichenauer S 2020 Anomaly detection with tensor networks (arXiv:2006.02516)
- [6] Liu J-G, Wang L and Zhang P 2021 Tropical tensor network for ground states of spin glasses *Phys. Rev. Lett.* **126** 090506

- [7] Liu J-G, Gao X, Cain M, Lukin M D and Wang S-T 2022 Computing solution space properties of combinatorial optimization problems via generic tensor networks (arXiv:2205.03718)
- [8] Hao T, Huang X, Jia C and Peng C 2022 A quantum-inspired tensor network method for constrained combinatorial optimization problems (arXiv:2203.15246)
- [9] Pozas-Kerstjens A, Hernández-Santana S, Monturiol J P, López M C, Scarpa G, González-Guillén C E and Pérez-García D 2022 Physics solutions for machine learning privacy leaks (arXiv:2202.12319)
- [10] Hornik K, Stinchcombe M and White H 1989 Multilayer feedforward networks are universal approximators *Neural Netw.* **2** 359–66
- [11] Glasser I, Sweke R, Pancotti N, Eisert J and Cirac I 2019 Expressive power of tensor-network factorizations for probabilistic modeling *Advances in Neural Information Processing Systems* vol 32
- [12] Markowitz H 1952 Portfolio selection *J. Finance* **7** 77–91
- [13] Conforti M, Cornuéjols G and Zambelli G 2014 *Integer Programming* vol 271 (Berlin: Springer)
- [14] Gleixner A et al 2021 Miplib 2017: data-driven compilation of the 6th mixed-integer programming library *Math. Program. Comput.* **13** 443–90
- [15] LLC Gurobi Optimization 2018 Gurobi optimizer reference manual (available at: www.gurobi.com/documentation/9.5/refman/index.html)
- [16] Hauke P, Katzgraber H G, Lechner W, Nishimori H and Oliver W D 2020 Perspectives of quantum annealing: methods and implementations *Rep. Prog. Phys.* **83** 054401
- [17] Nair V et al 2020 Solving mixed integer programs using neural networks (arXiv:2012.13349)
- [18] Cappart Q, Chételat D, Khalil E, Lodi A, Morris C and Veličković P 2021 Combinatorial optimization and reasoning with graph neural networks (arXiv:2102.09544)
- [19] Schuetz M J A, Brubaker J K and Katzgraber H G 2022 Combinatorial optimization with physics-inspired graph neural networks *Nat. Mach. Intell.* **4** 367–77
- [20] Gutmann S Farhi E, Goldstone J 2014 A quantum approximate optimization algorithm (arXiv:1411.4028)
- [21] Alcazar J, Vakili M G, Kalayci C B and Perdomo-Ortiz A 2021 Geo: enhancing combinatorial optimization with classical and quantum generative models (arXiv:2101.06250)
- [22] Schrijver A 1998 *Theory of Linear and Integer Programming* (New York: Wiley)
- [23] White S R 1992 Density matrix formulation for quantum renormalization groups *Phys. Rev. Lett.* **69** 2863
- [24] Schollwöck U 2011 The density-matrix renormalization group in the age of matrix product states *Ann. Phys., NY* **326** 96–192
- [25] Ferris A J and Vidal G 2012 Perfect sampling with unitary tensor networks *Phys. Rev. B* **85** 165146
- [26] In particular, a large class of optimization problems contain inequality type constraints and even nonlinear constraints
- [27] Glover F, Kochenberger G and Du Y 2018 A tutorial on formulating and using qubo models (arXiv:1811.11538)
- [28] Hibat-Allah M, Inack E M, Wiersema R, Melko R G and Carrasquilla J 2021 Variational neural annealing *Nat. Mach. Intell.* **3** 952–61
- [29] Bengio E, Jain M, Korablyov M, Precup D and Bengio Y 2021 Flow network based generative models for non-iterative diverse candidate generation (arXiv:2106.04399)
- [30] Hastings M B 2006 Solving gapped Hamiltonians locally *Phys. Rev. B* **73** 085115
- [31] Wolf M M, Verstraete F, Hastings M B and Cirac J I 2008 Area laws in quantum systems: mutual information and correlations *Phys. Rev. Lett.* **100** 070502
- [32] Orús R 2014 A practical introduction to tensor networks: matrix product states and projected entangled pair states *Ann. Phys., NY* **349** 117–58
- [33] Singh S, Pfeifer R N C and Vidal G 2010 Tensor network decompositions in the presence of a global symmetry *Phys. Rev. A* **82** 050301
- [34] Singh S, Pfeifer R N C and Vidal G 2011 Tensor network states and algorithms in the presence of a global U(1) symmetry *Phys. Rev. B* **83** 115125
- [35] Kleinberg J and Tardos E 2006 *Algorithm Design* (Boston: Pearson Education India) (available at: <https://ict.iitk.ac.in/wp-content/uploads/CS345-Algorithms-II-Algorithm-Design-by-Jon-Kleinberg-Eva-Tardos.pdf>)
- [36] Horowitz E and Sahni S 1974 Computing partitions with applications to the knapsack problem *J. ACM* **21** 277–92
- [37] Bellman R J N J 1957 *Dynamic Programming* (Princeton, NJ: Princeton University Press)
- [38] Fishman M, White S R and Stoudenmire E M 2020 The itensor software library for tensor network calculations (arXiv:2007.14822)
- [39] Sachdev S 1999 Quantum phase transitions *Phys. World* **12** 33
- [40] Gili K, Mauri M and Perdomo-Ortiz A 2022 Evaluating generalization in classical and quantum generative models (arXiv:2201.08770)
- [41] Vidal G 2007 Entanglement renormalization *Phys. Rev. Lett.* **99** 220405
- [42] Hastie T, Tibshirani R, Friedman J H and Friedman J H 2009 *The Elements of Statistical Learning: Data Mining, Inference and Prediction* vol 2 (New York: Springer)
- [43] Strashko A and Stoudenmire E M 2022 Generalization and overfitting in matrix product state machine learning architectures (arXiv:2208.04372)
- [44] Garey M R and Johnson D S 1979 *Computers and Intractability. A Guide to the Theory of NP-Completeness* (New York: W.H. Freeman and Co.)
- [45] Kellerer H, Mansini R and Speranza M G 2000 Selecting portfolios with fixed costs and minimum transaction lots *Ann. Oper. Res.* **99** 287–304
- [46] Deng D-L, Li X and Sarma S D 2017 Machine learning topological states *Phys. Rev. B* **96** 195145
- [47] Huang Y and Moore J E 2021 Neural network representation of tensor network and chiral states *Phys. Rev. Lett.* **127** 170601
- [48] Choo K, Carleo G, Regnault N and Neupert T 2018 Symmetries and many-body excitations with neural-network quantum states *Phys. Rev. Lett.* **121** 167204
- [49] Dawid A et al 2022 Modern applications of machine learning in quantum sciences (arXiv:2204.04198)
- [50] Shalev-Shwartz S, Shamir O and Shammah S 2017 Failures of gradient-based deep learning *Int. Conf. on Machine Learning* (PMLR) pp 3067–75
- [51] Romero E, Mazzanti F, Delgado J and Buchaca D 2019 Weighted contrastive divergence *Neural Netw.* **114** 147–56
- [52] Montúfar G F, Rauh J and Ay N 2011 Expressive power and approximation errors of restricted Boltzmann machines *Advances in Neural Information Processing Systems* vol 24
- [53] Stokes J and Terilla J 2019 Probabilistic modeling with matrix product states *Entropy* **21** 1236
- [54] Bradley T-D, Stoudenmire E M and Terilla J 2020 Modeling sequences with quantum states: a look under the hood *Mach. Learn.: Sci. Technol.* **1** 035008

- [55] McCulloch I P 2007 From density-matrix renormalization group to matrix product states *J. Stat. Mech.* **2007** 10014
- [56] Singh S and Vidal G 2012 Tensor network states and algorithms in the presence of a global SU(2) symmetry *Phys. Rev. B* **86** 195114
- [57] Singh S, Zhou H-Q and Vidal G 2010 Simulation of one-dimensional quantum systems with a global SU(2) symmetry *New J. Phys.* **12** 033029
- [58] Gasse M, Chételat D, Ferroni N, Charlin L and Lodi A 2019 Exact combinatorial optimization with graph convolutional neural networks *Advances in Neural Information Processing Systems* vol 32
- [59] Wang Z, Hadfield S, Jiang Z and Rieffel E G 2018 Quantum approximate optimization algorithm for maxcut: a fermionic view *Phys. Rev. A* **97** 022304
- [60] Niroula P, Shaydulin R, Yalovetzky R, Minssen P, Herman D, Hu S and Pistoia M 2022 Constrained quantum optimization for extractive summarization on a trapped-ion quantum computer *Sci. Rep.* **12** 1–14
- [61] Kondor R and Trivedi S 2018 On the generalization of equivariance and convolution in neural networks to the action of compact groups *Int. Conf. on Machine Learning* (PMLR) pp 2747–55
- [62] Bronstein M M, Bruna J, Cohen T and Velicković P 2021 Geometric deep learning: grids, groups, graphs, geodesics, and gauges (arXiv:2104.13478)
- [63] Meyer J J, Mularski M, Gil-Fuster E, Mele A A, Arzani F, Wilms A and Eisert J 2022 Exploiting symmetry in variational quantum machine learning (arXiv:2205.06217)
- [64] Larocca M, Sauvage F, Sbahi F M, Verdon G, Coles P J and Cerezo M 2022 Group-invariant quantum machine learning *PRX Quantum* **3** 030341
- [65] Nguyen Q T, Schatzki L, Braccia P, Ragone M, Coles P J, Sauvage F, Larocca M and Cerezo M 2022 Theory for equivariant quantum neural networks (arXiv:2210.08566)
- [66] Ragone M, Braccia P, Nguyen Q T, Schatzki L, Coles P J, Sauvage F, Larocca M and Cerezo M 2022 Representation theory for geometric quantum machine learning (arXiv:2210.07980)
- [67] Luo D, Chen Z, Hu K, Zhao Z, Hur V M and Clark B K 2021 Gauge invariant autoregressive neural networks for quantum lattice models (arXiv:2101.07243)
- [68] Luo D, Yuan S, Stokes J and Clark B K 2022 Gauge equivariant neural networks for 2+ 1d U(1) gauge theory simulations in Hamiltonian formulation (arXiv:2211.03198)
- [69] Hibat-Allah M, Ganahl M, Hayward L E, Melko R G and Carrasquilla J 2020 Recurrent neural network wave functions *Phys. Rev. Res.* **2** 023358
- [70] Morawetz S, De Vlught I J S, Carrasquilla J and Melko R G 2021 U(1)-symmetric recurrent neural networks for quantum state reconstruction *Phys. Rev. A* **104** 012401
- [71] Rudolph M S, Chen J, Miller J, Acharya A and Perdomo-Ortiz A 2022 Decomposition of matrix product states into shallow quantum circuits (arXiv:2209.00595)
- [72] Rudolph M S, Miller J, Chen J, Acharya A and Perdomo-Ortiz A 2022 Synergy between quantum circuits and tensor networks: short-cutting the race to practical quantum advantage (arXiv:2208.13673)
- [73] Benedetti M, Garcia-Pintos D, Perdomo O, Leyton-Ortega V, Nam Y and Perdomo-Ortiz A 2019 A generative modeling approach for benchmarking and training shallow quantum circuits *npj Quantum Inf.* **5** 45
- [74] Hemmecke R, Köppe M, Lee J and Weismantel R 2010 *Nonlinear Integer Programming* (Berlin: Springer)
- [75] Perez-Garcia D, Verstraete F, Wolf M M and Cirac J I 2007 Matrix product state representations *Quantum Inf. Comput.* **7** 401–30