

Received December 12, 2019, accepted March 6, 2020, date of publication April 7, 2020, date of current version April 22, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2986279

From Probabilistic Graphical Models to Generalized Tensor Networks for Supervised Learning

IVAN GLASSER^{ID}, NICOLA PANCOTTI^{ID}, AND J. IGNACIO CIRAC^{ID}

Max-Planck-Institut für Quantenoptik, 85748 Garching, Germany
Munich Center for Quantum Science and Technology (MCQST), 80799 Munich, Germany

Corresponding author: Ivan Glasser (ivan.glasser@mpq.mpg.de)

This work was supported in part by the ERC Advanced Grant QENOCOBA through the EU Horizon 2020 Program under Grant 742102, and in part by the German Research Foundation (DFG) through Germany's Excellence Strategy under Project EXC-2111 - 390814868 (MCQST). The work of Nicola Pancotti was supported by the ExQM.

ABSTRACT Tensor networks have found a wide use in a variety of applications in physics and computer science, recently leading to both theoretical insights as well as practical algorithms in machine learning. In this work we explore the connection between tensor networks and probabilistic graphical models, and show that it motivates the definition of generalized tensor networks where information from a tensor can be copied and reused in other parts of the network. We discuss the relationship between generalized tensor network architectures used in quantum physics, such as string-bond states, and architectures commonly used in machine learning. We provide an algorithm to train these networks in a supervised-learning context and show that they overcome the limitations of regular tensor networks in higher dimensions, while keeping the computation efficient. A method to combine neural networks and tensor networks as part of a common deep learning architecture is also introduced. We benchmark our algorithm for several generalized tensor network architectures on the task of classifying images and sounds, and show that they outperform previously introduced tensor-network algorithms. The models we consider also have a natural implementation on a quantum computer and may guide the development of near-term quantum machine learning architectures.

INDEX TERMS Boltzmann machines, graphical models, machine learning, quantum circuits, string-bond states, supervised learning, tensor networks, tensor-train.

I. INTRODUCTION

Tensor networks, which factorize a high-order tensor into a network of low-order tensors, have found a wide use of applications from quantum physics [1], [2] to machine learning [3], [4]. They can be used to compress weights of neural networks [5]–[8], to study model expressivity [9]–[14] or to parametrize complex dependencies between variables [15]–[18]. Recently, they have also attracted attention in the context of quantum machine learning [19], [20]: there has been much interest in understanding how low-depth quantum circuits that can be implemented on near-term quantum devices may be useful in machine learning [21]–[26], and tensor networks are a natural tool to perform the classical simulation of such algorithms [14], [27], [28]. Tensor networks that can be efficiently contracted on classical

computers thus provide a unique platform to benchmark and guide the development of new quantum machine learning architectures.

In this work, we explore the relationship between tensor networks and more common machine learning architectures, in particular probabilistic graphical models. We define generalized tensor networks which connect the two frameworks. These networks rely on the copy and reuse of local tensor information. Unlike regular tensor networks, they can be defined in complex geometries while remaining efficient to contract as long as an appropriate hierarchical order can be defined. We apply several variants of generalized tensor networks to image classification and environmental sound recognition and compare their performance, concluding that generalized tensor networks typically perform better than tensor networks alone.

Since they use several copies of the data inputs, generalized tensor networks share some structure with convolutional

The associate editor coordinating the review of this manuscript and approving it for publication was Chi-Yuan Chen^{ID}.

neural networks (CNN) [29], [30], while having direct connections to restricted Boltzmann machines (RBM) [31], [32]. Examples of such networks that have been used in quantum physics include string-bond states (SBS) [33], [34] (that generalize matrix product states (MPS)/tensor-trains [35]) as well as entangled plaquette states (EPS) [36]–[38]. Particular generalized tensor networks with tree structures have also been used to study the expressivity of deep learning models [10], [39].

We introduce an algorithm for performing supervised learning with generalized tensor networks, which combines stochastic gradient descent with previously introduced approaches for tensor networks. This framework generalizes works based on regular tensor networks such as MPS [16], [17], [40], [41] or tree tensor networks [42]. It has the advantage that more complex structures can be formed while keeping the computation efficient. This is especially useful for data that possesses some geometrical structure in more than one dimension, such as images. We emphasize that the algorithm does not need to rely on any Monte Carlo techniques, unlike in quantum physics, where generalized tensor networks can only be optimized in combination with computationally expensive Monte Carlo sampling. In particular the cost of optimizing a SBS is only a constant factor times the cost of optimizing a MPS, but SBS are much more flexible in higher dimensions and can interpolate between a MPS and a restricted Boltzmann machine.

We discuss how real-valued data can be used in conjunction with tensor networks and suggest to learn the relevant tensor features of real data as part of the network. Inspired by deep network architectures, we also propose two ideas to combine neural networks and tensor networks. In the first case we use a neural network to extract features from the data in order to feed them into a tensor network, in the second we combine generalized tensor networks and neural networks in the same deep network architecture.

We benchmark our algorithms for several generalized tensor network architectures on different data sets [43]–[45]. For image classification, we find that generalized tensor networks outperform previously introduced tensor-network algorithms based on MPS or trees while keeping a small dimension of the tensors. In the context of environmental sound recognition, we find that MPS and SBS deliver comparable performance. This shows that SBS should also be used along with MPS when considering one-dimensional data, especially in the presence of long-range correlations. Furthermore they may be applied in other settings such as natural language processing.

The network architectures we consider have a natural implementation on a quantum computer, which shows that simple quantum circuits that can be simulated classically can already achieve a good performance in supervised learning. It is still an open question whether quantum circuits that cannot be simulated classically will provide an advantage over classical machine-learning techniques. The networks we introduce may thus serve as a natural tool to test the performance and guide the development of such circuits.

II. GRAPHICAL MODELS AND GENERALIZED TENSOR NETWORKS

We first review definitions of probabilistic graphical models and tensor networks, discuss their relationship and show that the two frameworks can be connected through the definition of generalized tensor networks in which parts of the network can be copied and reused. Examples of generalized tensor networks which have been successfully used in quantum physics are introduced, and their connection to more common machine-learning architectures is discussed.

A. GRAPHICAL MODELS

Let us consider a set of discrete random variables $\mathbf{X} = \{X_1, \dots, X_N\}$ taking values $\mathbf{x} = (x_1, \dots, x_N)$ and a data set of samples from these variables $\mathcal{D} = \{\mathbf{d}_1, \dots, \mathbf{d}_{|D|}\}$. Inferring the underlying probability distribution $p(\mathbf{x})$ can be done by maximizing the log-likelihood

$$\mathcal{L} = \sum_{i=1}^{|D|} \log p(\mathbf{d}_i). \tag{1}$$

A common choice of parametrized models for p are graphical models [46], which correspond to a factorization of the probability distribution over a graph. Consider a graph $G = (V, E)$, where V is a set of vertices, E a set of edges between these vertices (each $e \in E$ is a pair of elements in V) and $\text{cl}(G)$ is the set of maximal cliques of the graph. An undirected graphical model or a Markov random field defines a factorization of the joint probability of all random variables as

$$p(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \prod_{C \in \text{cl}(G)} \phi_C(\mathbf{x}_C), \tag{2}$$

where \mathbf{x}_C are the values of the variables in clique C , ϕ_C are the clique potentials which are positive functions and Z is the partition function that ensures normalization of the probability (Fig. 1a). Graphical models can be converted to factor graphs [47] defined on a bipartite graph of factors and variable vertices: one factor node f_C is created for each maximal




FIGURE 1. (a) Undirected graphical model with three maximal cliques depicted in colors (b) Corresponding factor graph (c) Factor graph with hidden units in orange that are marginalized (d) Equivalent tensor network, which is a Matrix Product State.

clique and the factor is connected to the variables in the corresponding clique (Fig. 1b). The factorization of the probability distribution still reads

$$p(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \prod_C f_C(\mathbf{x}_C), \quad (3)$$

and inference can be performed through belief propagation and the sum-product algorithm on factor graphs. To increase the set of distributions which can be represented we can add additional dependencies by introducing ancillary hidden variables (which are unobserved, that is their values are not supplied in the data) $\mathbf{H} = \{H_1, \dots, H_M\}$ (Fig. 1c). The resulting probability distribution is obtained by marginalizing these hidden variables, giving

$$p(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \sum_{\mathbf{h}} \prod_C f_C(\mathbf{x}_C, \mathbf{h}_C). \quad (4)$$

B. TENSOR NETWORKS

Tensor networks are factorizations of tensors (multi-dimensional arrays) into a network of smaller tensors. This network admits a graphical notation, depicted in Fig. 2, in which boxes represent tensors and legs coming out of these boxes represent tensor indices. A leg joining two tensors represents a sum (contraction) over a joint index, such that the resulting tensor is a sum over joint indices of products of tensors. A simple example is a matrix factorization, in which a matrix is factorized in a product of two matrices. More formally, we consider a graph $G = (V, E)$ where we add to E edges containing only one vertex. These will correspond to open legs in a tensor-network graphical representation. We denote E' the subset of E containing edges that connect two vertices. We associate an integer size D_e called the bond dimension to each edge and define a tensor $T_v \in \otimes_{e \in \mathcal{V}_v} \mathbb{R}^{D_e}$ for each vertex $v \in V$, with indices associated with the edges of this vertex. A tensor-network state is defined by summing over indices on all edges in the graph that connect two tensors. The resulting contracted tensor network is a tensor indexed by

the indices of the open legs, denoted as $\mathbf{x} = (x_1, \dots, x_N)$:

$$T_{\mathbf{x}} = \sum_{e \in E'} \prod_v T_v. \quad (5)$$

A particular case are one-dimensional matrix product states (MPS) (Fig. 1d), also known as tensor-trains [35], which decompose a tensor as

$$T_{x_1, \dots, x_N} = \sum_{e \in E'} A_{e_1}^{x_1} A_{e_2, e_3}^{x_2} A_{e_3, e_4}^{x_3} \dots A_{e_N}^{x_N}, \quad (6)$$

where, for fixed value of \mathbf{x} , A^{x_1} and A^{x_N} are vectors, and A^{x_j} , $j = 2, \dots, N - 1$ are matrices. On a closed chain, also known as tensor ring [48], the corresponding decomposition is

$$T_{x_1, \dots, x_N} = \text{Tr} \left(\prod_{j=1, \dots, N} A^{x_j} \right), \quad (7)$$

where all A^{x_j} are matrices, and the dimension of these matrices is the bond dimension. Generalizations to trees and lattices in higher dimensions have also been studied.

C. TENSOR NETWORKS AND GRAPHICAL MODELS DUALITY

Consider a graphical model given as a factor graph in the special case where each hidden variable is connected to two factors, and each visible variable connected to one factor, as is the case in Fig. 1c. Because all the variables are discrete, each factor f_C is a non-negative function of discrete variables to which it is connected. It can therefore be written as a tensor $F_{C, \mathbf{x}_C, \mathbf{h}_C}$ with non-negative entries indexed by the variables \mathbf{x}_C and \mathbf{h}_C to which it is connected. The factor graph then defines a probability mass function over the visible variables that can be represented by a non-negative tensor $P_{\mathbf{x}}$ satisfying

$$P_{\mathbf{x}} = \sum_{\mathbf{h}} \prod_C F_{C, \mathbf{x}_C, \mathbf{h}_C}. \quad (8)$$

By comparing this expression with (5), we see that this probability mass function can be interpreted as a tensor-network state. Marginalization of a hidden variable corresponds to contracting the indices of the different factors connected to this variable, as indicated by the orange lines in Fig. 1d, and the visible variables correspond to the open legs of the tensor network. In the following we will refer to these open legs that correspond to visible variables or data inputs as the inputs of the tensor network. Probabilistic graphical models with discrete variables are therefore tensor networks of non-negative tensors. This connection has been previously observed in particular models [49]–[52], and [53] provides a more detailed analysis of this duality on hypergraphs. The fact that graphical models rely on non-negative tensors, whereas tensor networks are usually studied in the context of real (or complex) elements, has important consequences for the optimization algorithms. Graphical models can be used in conjunction with expectation-maximization algorithms, which rely on the computation of conditional probabilities over a subset of the variables. In general, tensor networks

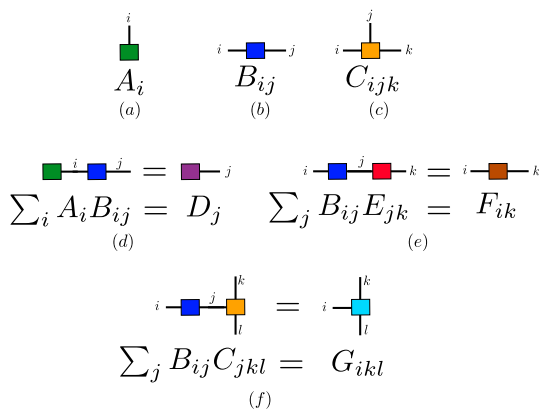


FIGURE 2. Graphical notation for tensor networks : (a) vector, (b) matrix, (c) order 3 tensor, (d) vector-matrix multiplication (e) matrix-matrix multiplication (f) matrix-tensor contraction.

cannot use the same algorithms, since they do not have the same probabilistic interpretation. However, they can use other powerful optimization algorithms, such as the density matrix renormalization group (DMRG). These algorithms can rely on the singular-value decomposition of real and complex matrices, without the constraint of non-negativity. A comparison of the expressive power of non-negative tensor networks and real or complex tensor networks in the context of probabilistic modeling has been performed in [14].

D. DUALITY WITH COPY TENSORS

Consider now the case of a probabilistic graphical model where one hidden variable is connected to several factors, as in Fig. 3a. To write the resulting probability mass function as a tensor network, we need to sum over the values of the hidden variable, which corresponds to summing over the value of the corresponding index which appears in several tensors. To represent this operation as a simple tensor network, we introduce a copy tensor δ with a number of indices corresponding to the number of factors connected to the hidden variable (each index can take as many values as the hidden variable), and such that δ is equal to one if all indices take the same value, and zero otherwise [54], [55]. In the case of Fig. 3a, δ is an order-3 tensor such that $\delta_{ijk} = 1$ if $i = j = k$ and $\delta_{ijk} = 0$ otherwise. We represent this tensor as a red dot in graphical representation. By inserting this copy tensor in the tensor network and connecting it to the corresponding factors, we obtain that the probability mass function defined by the graphical model can be represented by a tensor network, as in Fig. 3b.




FIGURE 3. (a) Factor graph with one hidden variable connected to several factors (b) Equivalent tensor network with copy tensor (c) Factor graph with visible variables connected to several factors (d) Equivalent tensor network (e) Corresponding tensor network if the visible variables have fixed input values.

The same procedure can be used if a visible variable is connected to several factors, as in Fig. 3c and Fig. 3d. Note that when computing the probability of inputs \mathbf{x} , the visible variables take specific given values. In this case contracting an input with a copy tensor will give the same result as fixing the value of all legs connected to the copy tensor to the input value and contracting the rest of the tensor network, as in Fig. 3e.

E. APPLICATION TO RESTRICTED BOLTZMANN MACHINES

One class of probabilistic graphical model which has the property that visible and hidden variables are connected to several factors are restricted Boltzmann machines (RBM) [31], [32]. They are defined on a bipartite graph with visible variables \mathbf{X} and hidden variables \mathbf{H} (Fig. 4a), so each visible variable is connected to each hidden variable. The connections between variables on this graph take the form of Ising interactions and the probability distribution of joint variables is

$$p(\mathbf{x}, \mathbf{h}) = \frac{1}{Z} e^{\mathcal{H}(\mathbf{x}, \mathbf{h})}, \tag{9}$$

where the Hamiltonian \mathcal{H} is a classical Ising Hamiltonian defined as (we omit here the bias terms for simplicity)

$$\mathcal{H} = \sum_{i,j} w_{ij} h_i x_j. \tag{10}$$

In the case where both visible and hidden variables are binary valued, the resulting probability distribution once the hidden variables have been marginalized is

$$p(\mathbf{x}) = \frac{1}{Z} \sum_{\mathbf{h}} e^{\mathcal{H}(\mathbf{x}, \mathbf{h})}, \tag{11}$$

$$= \frac{1}{Z} \prod_i (1 + e^{\sum_j w_{ij} x_j}). \tag{12}$$

Let us represent the corresponding factor graph, and associated tensor network. The cliques of the graph are all pairs of visible and hidden variables, for which there is an associated factor, which is a non-negative tensor, $f(x_j, h_i) = e^{w_{ij} h_i x_j}$.




FIGURE 4. (a) Restricted Boltzmann machine (RBM) (b) Corresponding tensor network (c) Corresponding tensor network when the input values are fixed and the input copy tensors contracted (d) Same tensor network with a reordering of the inputs and tensors.

By inserting a copy tensor at the position of the hidden variables, as well as of the input variables, we obtain the tensor network representation of the RBM in Fig. 4b [55].

F. RELATIONSHIP WITH MODELS USED IN QUANTUM PHYSICS

If we look at the graph for the tensor network equivalent to the restricted Boltzmann machine in Fig. 4b, we observe that this graph has many loops. Nevertheless, one can efficiently and exactly compute $p(\mathbf{x})$ for arbitrary size using (12). This comes from the fact that if we fix the values of the input variables and contract the copy tensors connected to the inputs, then the corresponding tensor network has no loops, as depicted in Fig. 4d.

Particular architectures of tensor networks satisfying the same property are tree tensor networks with copied inputs [10], [39]. Other examples of such networks have been used in the quantum physics community. The simplest example are entangled plaquette states (EPS) [36]–[38], also known as correlator product states, in which the tensor network is defined as a product of smaller tensors on overlapping clusters of variables:

$$T_{x_1, \dots, x_N} = \prod_{p=1}^P T_p^{x_p}, \tag{13}$$

where a coefficient $T_p^{x_p}$ is assigned to each of the 2^{n_p} (for binary variables) configurations \mathbf{x}_p of the variables in cluster p . Because the clusters overlap, the value of each variable is copied to all the tensors in which it is included (Fig. 5d). A sparse or short-range RBM is a special case of EPS [51], [52], in which the tensor $T_p^{x_p}$ takes the particular form $(1 + e^{\sum_j w_{pj}x_j})$, where the sum is limited to the variables in each cluster.

Another example are string-bond states (SBS) [33], [34], defined by placing MPS over strings (each string s is an ordered subset of the set of variables) on a graph which does not need to be a one-dimensional lattice. The resulting tensor network is

$$T_{x_1, \dots, x_N} = \prod_s \text{Tr} \left(\prod_{j \in s} A_{s,j}^{x_j} \right). \tag{14}$$

The value of each visible variable is copied and sent to different MPS (Fig. 5b). It was shown in [52] that a RBM is a special case of SBS for which each string is associated with a hidden variable and covers all visible variables, and the matrices are taken to be

$$A_{s,j}^{x_j} = \begin{pmatrix} 1 & 0 \\ 0 & e^{w_{sj}x_j} \end{pmatrix}. \tag{15}$$

SBS thus provide a generalization of RBM that is naturally defined for discrete variables of any dimension and can introduce different correlations through the use of higher dimensional and non-commuting matrices. Since SBS also include a MPS as a particular case, they provide a way to interpolate between a MPS (large bond dimension, only one string) and a RBM (bond dimension 2, diagonal matrices, many strings). Note that a product of S MPS of bond dimension D is an MPS with bond dimension D^S , so restricted Boltzmann machines can be written as MPS of large bond dimension but with diagonal matrices.

G. GENERALIZED TENSOR NETWORKS

1) INPUT FEATURES FOR TENSOR NETWORKS

In many cases of interest, data is not discrete, but instead given in the form of real numbers. In order to train tensor networks to perform machine learning tasks on such data, it has been suggested to define feature vectors from the data, and then to contract these feature vectors with a tensor network to define a function of the inputs [16], [17]. Consider for example input data given as real numbers $\{x_1, \dots, x_N\}$ and define feature vectors

$$v_i = \begin{pmatrix} 1 \\ x_i \end{pmatrix}. \tag{16}$$

Now it is possible to use a tensor network to define a function of the inputs by contracting these vectors with the open legs of a tensor network, as depicted in Fig. 6 in the case of a MPS.

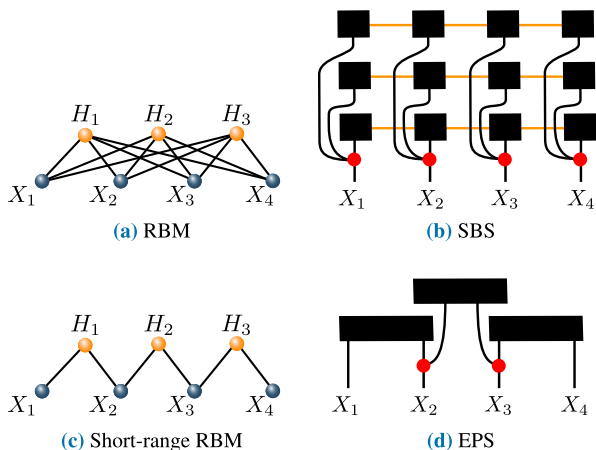


FIGURE 5. (a) Restricted Boltzmann machine (RBM) consisting of visible and hidden variables (b) String-bond state with 1D geometry generalizing RBM. The legs corresponding to contracted indices in each MPS are depicted in orange for visibility. (c) Short-range RBM with local connections between visible and hidden variables (d) Entangled plaquette state (EPS) generalizing the short-range RBM.




FIGURE 6. MPS with feature vectors as inputs.

2) COPY OPERATION WITH VECTOR INPUTS

We would like to extend the previous results that use copy tensors to copy the inputs of the tensor network, allowing to efficiently contract complex tensor networks when the inputs are given. This idea that input data may be copied is also used in other machine learning architectures, such as for example Convolutional Neural Networks (CNN): a convolutional layer applies several filters to local clusters of input variables. In order to apply different filters to the same input data, this data needs to be copied and fed to the different parts of the neural network. In the case of tensor networks, we face the problem that copy tensors can only be used to copy discrete inputs [39]. Placing copy tensors as the input of the tensor network will therefore not have the expected effect.

For these reasons we introduce a different notion of copy, that we call the copy operation, that allows to copy real vectors. We will call the models that use this feature generalized tensor networks, to distinguish them from tensor networks with copy tensors, and we will restrict ourselves to the special cases in which these networks can be efficiently contracted as long as the inputs have fixed values. In the special case where the copy only takes place at the level of the input variables, and that these variables are discrete, then these networks will coincide with tensor networks that use a copy tensor at the level of the input legs.

We graphically depict the copy operation through a red dot between edges of the graph and an arrow which marks the incoming edge. The copy operation takes a vector as input, and it outputs two identical copies of it (Fig. 7). This operation is equivalent to having two copies of the input contracted with the rest of the tensor network, and we depict this fact by using the same colors for the same tensors. In the case where the input is a discrete variable, we define the copy operation so that it has the same effect as the copy tensor that simply copies this variable.

More generally, we can use this property to define complex generalized tensor networks where the copy operation is used at different places in the network. The resulting tensor networks are well defined if they can be expressed as a tensor network using several times the same tensor. Generalized tensor networks are therefore tensor networks that use weight sharing between some tensors, and such that several copies of the inputs may be used. In Fig. 7 we provide examples of such a mapping to clarify the definition of generalized tensor networks.

3) ARCHITECTURES USED IN THIS WORK

The previous examples of EPS and SBS can be straightforwardly extended to generalized tensor networks by replacing the copy tensors as inputs by copy operations, so that they can be used with feature vector inputs. More generally, one can think of complex networks built using the copy operation for gluing different networks together. In practice, we will consider the following network architectures, in the case

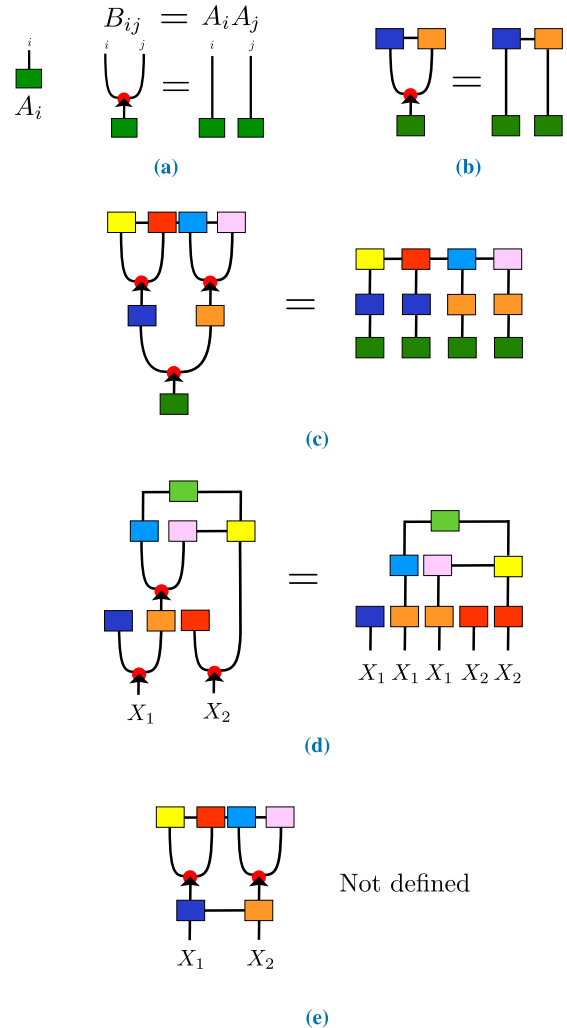


FIGURE 7. (a) Copy operation of a vector input A_j , resulting in a new tensor $B_{ij} = A_i A_j$ (b), (c) and (d) : Mapping of generalized tensor networks with copy operation to tensor networks with weight sharing (indicated by the same colors of the tensors) and copied input vectors or discrete inputs. (e) This representation is not defined because the copy operation only applies to vectors.

where the inputs have a two-dimensional geometry, such as images:

- EPS with 2×2 overlapping plaquettes with weight sharing such that the tensor for each plaquette is the same (Fig. 8a).
- SBS defined with horizontal and vertical strings covering the 2D lattice (Fig. 8b). We will denote this kind of SBS as 2D-SBS. Correlations along one of the two dimensions can be captured in the corresponding MPS, and more complex correlations are included through the overlap of the different strings.
- SBS consisting of 4 strings, each covering the whole lattice in a snake pattern, but in a different order (Fig. 8c). We denote these SBS as snake-SBS. They have the advantage, compared to a MPS, that two nearest neighbours variables always appear next to each other in one

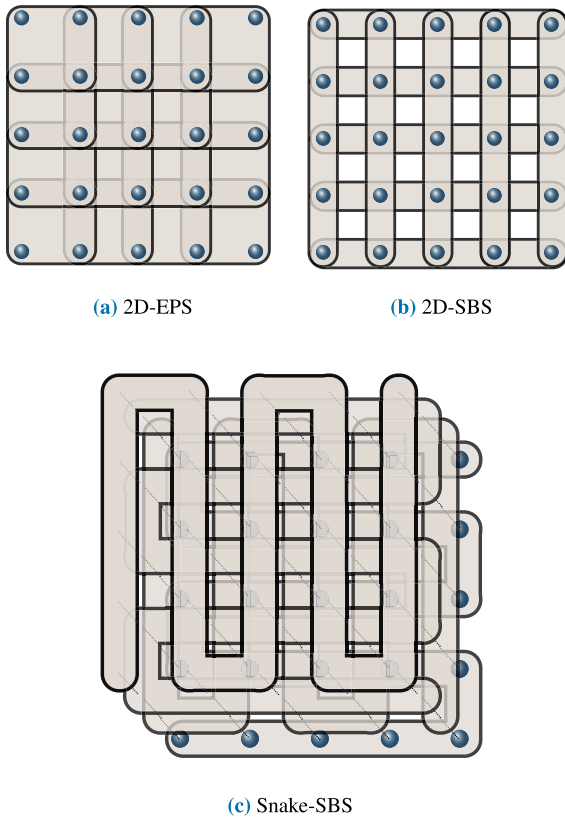


FIGURE 8. Possible geometries of generalized tensor networks in 2D: (a) EPS. (b) 2D-SBS consisting of horizontal and vertical overlapping strings. (c) Snake-SBS consisting of 4 overlapping strings in a snake pattern.

of the 4 strings, thus rendering the capture of strong local correlations efficient. More complex string geometries could be considered, and the choice of string could be itself learned with a RBM.

- EPS with an extra output leg in each plaquette, and such that the outputs are copied and taken as input into a SBS (Fig. 9). The input variables are first copied and fed into overlapping clusters parametrized by tensors. The output leg of each of these tensors is a vector which is copied a few times. Each of these copies can then be contracted with the open legs of a different MPS, forming together




FIGURE 9. EPS-SBS consisting of a first layer of EPS, followed by a copy operation and a second layer of SBS.

a SBS. In 2D, we choose 2×2 overlapping plaquettes in the first layer, and 4 strings forming a snake-SBS in the second layer. In the following we will call this generalized tensor network EPS-SBS.

These are simple examples of generalized tensor networks inspired from models invented in quantum physics to capture correlations on 2D lattice systems. More complex networks based on trees or hierarchical designs with more than two layers can also be constructed in the same way.

We note that there is a relationship between EPS and one layer of a CNN : in the case of discrete inputs, there is a finite number of possible different filters that are applied in a CNN. If we represent the local convolutional operation as a function from discrete inputs, then this function can be represented as a tensor. The fact that convolutional filters are applied with a small stride over an input image corresponds to the overlap of the plaquettes in an EPS, and so EPS correspond to the first layer of a CNN applying all possible filters on discrete inputs.

These generalized tensor networks have the advantage, compared to standard tensor networks, that they can be easily defined in arbitrary dimension and geometry while remaining efficient to contract, as long as the input is fixed. This is in contrast to a 2D tensor network previously introduced in physics known as projected entangled pair states [56], which is naturally defined in higher dimensions, but cannot be contracted exactly efficiently. In particular, 2D-SBS form a subclass of projected entangled pair states that remains efficient to contract.

Another advantage of generalized tensor networks is that they can represent some functions with fewer parameters than regular tensor networks without copy tensors. Indeed, tensor networks such as MPS or tree tensor networks satisfy a constraint known as area law, which implies that they will require a number of parameters that is exponential in the number of variables to represent functions not satisfying this constraint, whereas generalized tensor networks such as SBS or RBM can represent some of these functions with a polynomial number of parameter [39], [57].

III. SUPERVISED LEARNING ALGORITHM

Graphical models are often used in conjunction with unsupervised learning algorithms, since they are designed to represent probability distributions. In particular cases it is possible to compute the normalization Z , which gives exact access to the likelihood and makes maximum likelihood estimation tractable. This is possible for graphical models and tensor networks on trees and has led to an algorithm for generative modeling with MPS [14], [40]. In the more general case, which includes RBM, the normalization Z cannot be computed efficiently. Approximate algorithms relying on Monte Carlo sampling can then be used, such as contrastive divergence [32], [58]. Generalized tensor networks suffer from the same issue, which makes unsupervised learning computationally expensive. Since these networks correspond to quantum states, it might be possible to implement them on a quantum

computer and sample from them efficiently. In this work we focus instead on supervised learning, where access to the normalization Z is not necessary. In this section we discuss how RBM can be used for supervised learning, and generalize the corresponding algorithm to tensor networks.

A. SUPERVISED LEARNING WITH RESTRICTED BOLTZMANN MACHINES

We first review how RBM can be used to perform supervised learning, in a classification setting [59], [60]. Given labelled training data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$, where the y_i take discrete values corresponding to different classes, a RBM can be used to approximate the joint probability distribution of the variables and labels:

$$p(\mathbf{x}, y) = \frac{1}{Z} \sum_{\mathbf{h}} e^{\mathcal{H}(\mathbf{x}, \mathbf{h}, y)} \quad (17)$$

In such a model, the label is seen as an additional visible variable (Fig. 10a), possibly encoded in a one-hot representation to use only binary units. Training such a generative model can be done by maximizing the log-likelihood of the data. Since the likelihood is intractable, because the partition function Z cannot be efficiently computed, such a training can be done through approximate algorithms such as contrastive divergence. In supervised learning, one is interested in computing the conditional distribution

$$p(y|\mathbf{x}) = \frac{p(\mathbf{x}, y)}{\sum_{y_j} p(\mathbf{x}, y_j)}, \quad (18)$$

which can be computed exactly when the number of classes is small enough, since the two partition functions cancel. The label predicted by the model for new data \mathbf{x}_i is the label maximizing $p(y|\mathbf{x}_i)$. Since one is ultimately interested in classification performance, it can be advantageous to directly optimize $p(y_i|\mathbf{x}_i)$, which leads to a cost function to minimize

$$\mathcal{L}_{\text{discriminative}} = - \sum_{i=1}^{|\mathcal{D}|} \log p(y_i|\mathbf{x}_i), \quad (19)$$

whose gradient can be computed analytically. A RBM can moreover be trained in a semi-supervised way, by using a combination of discriminative and generative training.

B. SUPERVISED LEARNING WITH GENERALIZED TENSOR NETWORKS

To generalize the discriminative training of RBM to generalized tensor networks, we would like to approximate the joint probability distribution of the variables and labels as a tensor network:

$$p(\mathbf{x}, y) \propto \text{GTN}(\mathbf{x}, y), \quad (20)$$

where $\text{GTN}(\mathbf{x}, y)$ is the function resulting of the contraction of a generalized tensor network with the inputs features and with the discrete label variable. More specifically, in the case of discrete variables the variables \mathbf{x} and y fix the values of the open legs of the generalized tensor network, and the whole






FIGURE 10. (a) A classification RBM turns the label into an additional visible unit. (b) The same procedure can be defined for a SBS by adding a node corresponding to the label, and corresponding tensors which connect it to the rest of the tensor network. (c) Generalized tensor networks can be combined with additional layers of neural networks. For example an EPS output is a tensor that can be combined with a linear classifier.

network is contracted, resulting in a scalar. In the case of input variables that are feature vectors, each vector is contracted with the corresponding open leg of the tensor network. In order for the cost function to be well defined, the result of the network contraction should be positive. This can be done by ensuring that the tensor elements are positive, as in a RBM or graphical model, or by choosing instead $p(\mathbf{x}, y) \propto (\text{GTN}(\mathbf{x}, y))^2$, which corresponds to a Born machine [14], [61], [62], or $p(\mathbf{x}, y) \propto e^{\text{GTN}(\mathbf{x}, y)}$. In the following we will adopt this last choice, for which training is found to be more efficient. The label is now seen as the index of one tensor. Since it is discrete, there is no need to use a one-hot representation and one can simply enlarge the dimension of the leg of a tensor to accommodate for the number of possible classes (Fig. 10b). We then define, by analogy with the graphical model case,

$$p(y_k|\mathbf{x}_i) = \frac{e^{\text{GTN}(\mathbf{x}_i, y_k)}}{\sum_{y_j} e^{\text{GTN}(\mathbf{x}_i, y_j)}}, \quad (21)$$

which can be seen as a softmax function applied to the different outputs of the tensor network contraction, and the cost function is again chosen to be a cross-entropy loss

$$\mathcal{L}_{\text{discriminative}} = - \sum_{i=1}^{|\mathcal{D}|} \log p(y_i|\mathbf{x}_i). \quad (22)$$

The gradient of the cost function can then be expressed as a function of $\text{GTN}(\mathbf{x}_i, y_j)$ and its derivatives, since

$$\frac{\partial \log p(y_i|\mathbf{x}_i)}{\partial w} = \frac{\partial \text{GTN}(\mathbf{x}_i, y_i)}{\partial w} - \sum_{y_j} p(y_j|\mathbf{x}_i) \frac{\partial \text{GTN}(\mathbf{x}_i, y_j)}{\partial w}. \tag{23}$$

$\text{GTN}(\mathbf{x}_i, y_j)$ can be computed exactly by fixing the value of the input units and labels and contracting the network. We note that in general we can contract the whole network without the labels, and perform the contraction for the different labels as a last step. Contraction of the whole network for different labels thus only adds a small cost (which depends on the shape of the network) to the contraction of the network without labels. We observe that from the point of view of supervised learning there is no essential difference between SBS and MPS in terms of the optimization algorithms: the cost of optimizing a SBS is only a constant factor (the number of strings) more expensive than for a MPS. This is unlike in quantum physics where Monte Carlo sampling is necessary to optimize a SBS. To further regularize the tensor network, we randomly drop tensor elements to 0 with probability δ during training, following the procedure in [17].

So far we have constructed tensor networks which, when an input and a label is given, have no open legs. We can also construct networks with open legs and use tensor networks in combination with other machine learning techniques. In this case the tensor network maps the input to a tensor which can for example be used as input in a neural network. In the case of EPS where each tensor over overlapping plaquettes has an open leg, an input is mapped to a tensor with an extra dimension as output. The simplest way to combine EPS with other neural networks is to place a linear classifier on top of the EPS (Fig. 10c). The backpropagation algorithm used to compute derivatives of the neural network is in this case combined with the algorithm for computing derivatives of a tensor network, and the joint network can be optimized using stochastic gradient descent. More complex combinations of tensor networks and neural networks may be defined in the same way.

IV. LEARNING FEATURE VECTORS OF DATA

In this section we explore several strategies that can be used to deal with data that is not discrete. We suggest to learn relevant tensor features as part of the tensor network and discuss how tensor features can also be learned as part of a deep learning architectures which combines a neural network extracting features with a tensor network.

A naive way of applying tensor networks with real data would be to discretize data or use its binary representation. This is not a suitable approach, because that would amount to dramatically increasing the size of the data, rendering learning very slow, and it would also lead to large tensor networks prone to overfitting. Another approach, as suggested by [16] and [17], is to map the real data to a higher dimensional feature space. Each variable is first independently




FIGURE 11. (a) Real inputs X_i are mapped to a feature vector (here with length two). This vector can then be used as input to a generalized tensor network by contracting it with the open legs of the generalized tensor network. (b) Feature tensors can compress the discretized representation of the inputs X_i to a smaller dimensional space. These tensors can share weights and can be learned as part of the tensor network.

mapped to a vector of length (at least) two in order to be contracted with the open legs of the tensor network (Fig. 11a). Choices of feature maps that have been used include

$$x \rightarrow \begin{pmatrix} 1 \\ x \end{pmatrix} \text{ or } \begin{pmatrix} \cos(\frac{\pi}{2}x) \\ \sin(\frac{\pi}{2}x) \end{pmatrix}, \tag{24}$$

and generalizations to higher dimensions. A choice which is suitable with our algorithm, assuming that the data is normalized between 0 and 1, is to use

$$x \rightarrow \begin{pmatrix} \cos^2(\frac{\pi}{2}x) \\ \sin^2(\frac{\pi}{2}x) \end{pmatrix}, \tag{25}$$

because this ensures that the vectors are positive and the normalization prevents numerical instabilities.

These choices however put severe limitations to the functions that can be learned. Indeed, the data set with just two variables presented in Fig. 12a cannot be separated by a MPS of bond dimension 2 with one of these feature choices, since the boundary decision will be a polynomial of degree two of the features. Nevertheless, a different feature choice could distinguish the two classes, even with bond dimension 2. We therefore suggest to learn the appropriate features as part of the learning algorithm. This can be done by parametrizing the feature functions and learning them at the same time as the rest of the network. To be able to use a purely tensor network algorithm, we can parametrize these functions using a tensor network. In the simplest case, we discretize the real data and use a tensor to compress the large dimensional input into a smaller dimensional vector of suitable length, as shown in Fig. 11b. This tensor can be learned as part of the whole tensor network and prevents the size of the rest of the tensor network to increase when the discretization size changes. The feature tensor can be the same for all variables, for example image pixels, but can be different in the case where the variables are of different nature. Using this procedure, a MPS of bond dimension 2 is able to get perfect accuracy on the data set presented in Fig. 12a. Here the two variables X_1 and X_2 are continuous between 0 and 1. Each of these variables is discretized into a variable that can only take 16 different values between 0 and 1. A tensor T of size




FIGURE 12. (a) Dataset with two features X_1 and X_2 and two classes (depicted in different colors) that cannot be learned by a MPS of bond dimension 2 with the feature choice in (25). (b) Two normalized features learned by a tensor while classifying the previous data set with a MPS of bond dimension 2. The features have been discretized in 16 intervals. Using this choice of features the MPS can classify the data set perfectly.

16×2 then represents the tensor feature map. We represent in Fig. 12b the result T_{X_i} of the contraction between the discrete representation of X_i and the feature tensor T . This represents the two features that the network has learned, and corresponds to the functions that could have been used in the feature map in (24). We note that starting from random features on more complex data sets makes learning difficult, but the feature tensor can be pretrained using a linear classifier, before being trained with the rest of the network.




FIGURE 13. (a) Choice of two features in (25) for an input taking real values between 0 and 1. (b) Two normalized features learned by a tensor with output dimension 2 combined with a snake-SBS classifying the MNIST data set. The input features x are the greyscale value of pixels, normalized between 0 and 1 and discretized in 16 intervals.

In comparison, we also show in Fig. 13b the features learned while classifying MNIST with greyscale pixels and a snake-SBS (see Section V). These features are not very different from the choice in 25, and we could not distinguish performance with this choice or with learned features on this data set. We expect however that this procedure will be necessary for more complex data sets which are not easily approximated by a binary function. Moreover the size of the feature vector provides a regularization of the model, and higher sizes might be necessary for more complex data sets. More generally this tensor could be itself represented with a small tensor network, to prevent the number of parameters to increase too much with a very small discretization interval. It is interesting to note that the features learned in our examples are almost continuous even if we use smaller discretization intervals. This means that two real inputs that are close to each other will lead to the same predictions by the network, a property which is in general not true if we simply discretize the inputs and use a larger tensor network. Our approach of learning

the features as part of the tensor network may be especially relevant in the context of quantum machine learning, where the tensor network is replaced by a quantum circuit and it might be suitable to have the full network as part of the same quantum machine learning architecture.

As an alternative way of choosing the features, we can combine the feature choice with other machine learning techniques. If the input data represents images, it is a natural choice to use Convolutional Neural Networks as feature extractors, since these have been highly successful for image classification. CNN consist in convolution filters, which use convolutional kernels to transform an image into a set of filtered images, and pooling layers which downsize the images (Fig. 14). The resulting features preserve a form of locality. Therefore it is natural to consider the vector of applied filters associated with each location in the image as a feature vector that can be used in conjunction with generalized tensor networks. The CNN and the tensor network can be trained together, since the derivatives of the tensor network can be used in the backpropagation algorithm which computes the gradient of the cost function.




FIGURE 14. Using Convolutional Neural Networks as feature vector extractors from real data: the output of the CNN is seen as an image with a third dimension collecting the different features. For each pixel of this image, the vector of features is contracted with the open legs of a tensor network.

V. NUMERICAL EXPERIMENTS

We test the generalized tensor network approach on the task of image classification, where a natural two-dimensional geometry that can be reflected in the architecture of the tensor network is present, as well as on the task of urban sound recognition, where the time dimension provides a one-dimensional geometry.

A. IMAGE CLASSIFICATION

We first consider the MNIST data set [44], which consists of 28×28 greyscale images of digits. There are 10 classes and we adopt a multiclass classification procedure in which one tensor of the tensor network is parametrized by the ten possible labels. The original training set is split into training and validation sets of 55000 and 5000 examples and the performance of the different models is evaluated on the test set of 10000 examples. We consider the following generalized tensor networks: a snake-SBS with 4 strings (Fig. 8c), a 2D-SBS (Fig. 8b), an EPS with a 2×2 translational-invariant plaquette combined with a linear classifier, (Fig. 10c), an EPS-SBS

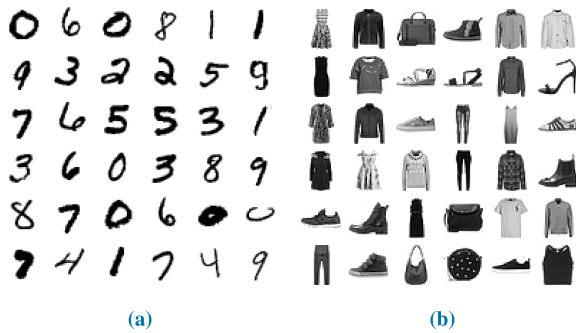


FIGURE 15. Examples of images from the MNIST (a) and fashion MNIST (b) data set.

with translational-invariant plaquette combined with a snake-SBS (Fig. 9) and a CNN-snake-SBS which uses a 1-layer CNN as input features (Fig. 14). The CNN considered here uses a convolutional layer applying 65×5 filters (stride 1) with ReLU activation function and a pooling layer performing max pooling with a 2×2 filter. All other networks use the choice of features presented in (25) and the greyscale values are normalized between 0 and 1. We compare the performance of these networks with a MPS [16] and a RBM (the number of hidden units of 250, 500, 750 or 1000 is taken as a hyperparameter). All networks use a batch size of 20 examples and hyperparameters such as the learning rate α and the regularization rate δ are determined through a grid search while evaluating the performance on the validation set. Best performance is typically achieved with $\alpha = 10^{-4}$, $\delta = 0.95$ and a hundred epochs of training.




FIGURE 16. Test set accuracy of different generalized tensor networks on the MNIST data set.

The test set accuracy, presented in Fig. 16, shows that even with a very small bond dimension generalized tensor network are able to accurately classify the data set. Their performance is significantly better than that of a tree tensor network [42] or a MPS trained in frequency space [41], and while a MPS can also achieve 99.03% accuracy with a bond dimension of 120 [16], the cost of optimizing very large tensors has prohibited the use of this method for larger problems so far. The snake-SBS with bond dimension larger than 6 has also better performance than a RBM. Since the snake-SBS provides an interpolation between RBM and MPS, the choice

of number of strings and geometry can be seen as additional parameters which could be tuned further to improve over the performance of both methods. All networks have a training set accuracy very close to 100% when the bond dimension is larger than 6, and we expect that better regularization techniques or network architectures have to be developed to significantly increase the test set performances obtained here. We also optimized a snake-SBS with positive elements (by parametrizing each element in a tensor as the exponential of the new parameters), which is a graphical model. Using the same algorithm, we were not able to achieve better performance than 93% classification accuracy with bond dimensions up to 10. This shows that while having a structure closely related to graphical models, tensor networks may provide different advantages.

TABLE 1. Test set accuracy of generalized tensor networks and other approaches [43] on the fashion MNIST data set.

Method	Accuracy
Support Vector Machine	84.1%
EPS + linear classifier	86.3%
Multilayer perceptron	87.7%
EPS-SBS	88.6%
Snake-SBS	89.2%
AlexNet	89.9%
CNN-snake-SBS	92.3%
GoogLeNet	93.7%

We then turn to the fashion MNIST data set [43], consisting of 28×28 greyscale images of clothes. While having the same size as the original MNIST data set, it is significantly harder to classify. We report the best accuracy obtained with different generalized tensor networks with bond dimension up to 10 in Table 1. It is found that these networks, while not state-of-the-art, are competitive with other approaches such as Support Vector Machines, AlexNet and GoogLeNet Convolutional Neural Networks or a multilayer perceptron neural network, which is encouraging considering the potential improvements in terms of network architecture or training algorithms. Advantages in favor of the use of tensor networks are the adaptive choice of hyperparameters such as the bond dimension [16], as well as the good theoretical understanding of their properties [2], but improvements are still needed before practical use in large-scale supervised learning.

B. ENVIRONMENTAL SOUND CLASSIFICATION

So far we have considered black and white images, but it is also interesting to study how generalized tensor networks could be used for other types of data. In the following we consider the task of classifying environmental sounds. The UrbanSound8K data set [63] is a collection of 8732 audio clips (4s or less) divided into 10 classes of urban sounds: air conditioner, car horn, children playing, dog barking, drilling, engine idling, gun shot, jackhammer, siren and street music. The data set is divided into 10 folds and we use folds 1-9

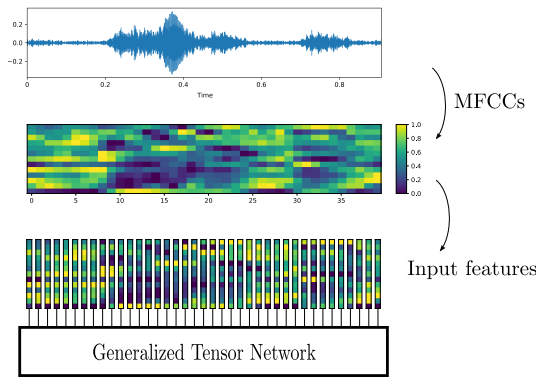


FIGURE 17. From the raw audio signal, Mel-frequency cepstral coefficients (MFCCs) are extracted over short overlapping windows, resulting in a sequence of high dimensional vectors. These vectors are taken as input to a generalized tensor network.




FIGURE 18. Training and testing accuracy of a MPS and a SBS with 4 strings on the UrbanSound8K data set. The density of parameters is the total number of parameters divided by 174 (the length of the strings).

for training and fold 10 for testing. The one-dimensional structure of sounds allows us to compare MPS and SBS with the same 1D string geometry. Preprocessing of the data takes place as follows : clips shorter than 4s are repeated to reach a fixed length of 4s. The first 13 Mel-frequency cepstral coefficients (MFCCs) are extracted for each clip (sampled at 22050Hz) using a window size of 2048 and hop length of 512, resulting in a sequence of length 173 and dimension 13 (Fig. 17). The corresponding 13-dimensional vectors are used as input feature vectors for the tensor network, and the time dimension of the sequence corresponds to the 1-dimensional structure of the MPS, or the strings of the SBS. Note that we do not perform any data augmentation nor split the training examples to enlarge the size of the data set, since we are interested in comparing MPS and SBS, rather than achieving the best possible accuracy on this data set. The training and testing accuracies are reported in Fig. 18 for a MPS with bond dimension up to 10 and a SBS with 4 strings and bond dimension up to 5. Since we are interested in comparing the expressivity of the different networks, no regularization is used and training is performed until the training accuracy does not improve anymore. Note that a MPS with bond dimension D has as many variational parameters as a SBS with 4 strings and bond dimension $D/2$.

We observe that the SBS has slightly higher training accuracy than a MPS with larger bond dimension and the

same number of parameters. The test set performance is not significantly different between distinct architectures and in both cases we find that a lot of overfitting has taken place, which is not surprising given the small number of training examples. Higher accuracies have been reported with other methods on the same data set. For example Convolutional Neural Networks can reach above 70% test set accuracy [45], however they use many more input features and rely on data augmentation. Nevertheless our results show that SBS should also be considered along with MPS when considering one-dimensional data, and may be applied in other settings such as natural language processing [64], [65].

VI. CONCLUSION

We have introduced generalized tensor networks, which enlarge the class of tensor networks by introducing a reuse of information taking the form of a copy operation of tensor elements. The resulting networks are related to graphical models and we have discussed the strong relations that exist between particular graphical models and tensor network structures, such as restricted Boltzmann machines and string-bond states. We provided an algorithm to train these models to perform a supervised learning task and discussed several strategies to use tensor networks in conjunction with real-valued data. We showed that generalized tensor networks that can be contracted exactly can perform accurate image classification with much smaller bond dimension than regular tensor networks, that they can be used in other settings such as sound recognition and that they can be combined with neural-network architectures. Tensor networks can also be seen as a tool to simulate quantum circuits, and there is much research trying to understand how quantum circuits can be used in machine learning. Quantum circuits corresponding to MPS or tree tensor networks have been studied in the context of quantum machine learning [16], [27]. To implement the function corresponding to a generalized tensor network, one would need to copy the data at the input of the quantum circuit. Our results thus show that quantum machine learning circuits should take as input several copies of each data input and not just a single one, as well as share parameters between the different unitaries in the circuit. Generalized tensor networks which originate from the classical simulation of quantum states may thus serve as a testing and benchmarking platform of near-term quantum machine learning architectures.

ACKNOWLEDGMENT

We would like to thank Shi-Ju Ran, Yoav Levine, and Miles Stoudenmire for discussions, and Vedran Dunjko for a careful reading of the paper and helpful comments.

REFERENCES


- [1] R. Orús, “A practical introduction to tensor networks: Matrix product states and projected entangled pair states,” *Ann. Phys.*, vol. 349, pp. 117–158, Oct. 2014.
- [2] F. Verstraete, V. Murg, and J. I. Cirac, “Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems,” *Adv. Phys.*, vol. 57, no. 2, pp. 143–224, Mar. 2008.

- [3] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, and D. P. Mandic, "Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions," *Found. Trends Mach. Learn.*, vol. 9, nos. 4–5, pp. 249–429, 2016.
- [4] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, M. Sugiyama, and D. P. Mandic, "Tensor networks for dimensionality reduction and large-scale optimization: Part 2 applications and future perspectives," *Found. Trends Mach. Learn.*, vol. 9, no. 6, pp. 249–429, 2017.
- [5] A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov, "Tensorizing neural networks," in *Proc. Annu. Conf. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2015, pp. 442–450.
- [6] A. Hallam, E. Grant, V. Stojevic, S. Severini, and A. G. Green, "Compact neural networks based on the multiscale entanglement renormalization Ansatz," 2017, *arXiv:1711.03357*. [Online]. Available: <http://arxiv.org/abs/1711.03357>
- [7] R. Yu, S. Zheng, A. Anandkumar, and Y. Yue, "Long-term forecasting using higher order tensor RNNs," 2017, *arXiv:1711.00073*. [Online]. Available: <http://arxiv.org/abs/1711.00073>
- [8] J. Kossaiji, Z. C. Lipton, A. Khanna, T. Furlanello, and A. Anandkumar, "Tensor regression networks," 2017, *arXiv:1707.08308*. [Online]. Available: <http://arxiv.org/abs/1707.08308>
- [9] N. Cohen, O. Sharir, Y. Levine, R. Tamari, D. Yakira, and A. Shashua, "On the expressive power of deep learning: A tensor analysis," in *Proc. 29th Conf. Learn. Theory (COLT)*, New York, NY, USA, Jun. 2016, pp. 698–728.
- [10] N. Cohen and A. Shashua, "Convolutional rectifier networks as generalized tensor decompositions," in *Proc. 33rd Int. Conf. Mach. Learn. (ICML)*, New York City, NY, USA, Jun. 2016, pp. 955–963.
- [11] N. Cohen and A. Shashua, "Inductive bias of deep convolutional networks through pooling geometry," in *Proc. 5th Int. Conf. Learn. Represent. (ICLR)*, Toulon, France, Apr. 2017, pp. 1–28.
- [12] N. Cohen, O. Sharir, Y. Levine, R. Tamari, D. Yakira, and A. Shashua, "Analysis and design of convolutional networks via hierarchical tensor decompositions," 2017, *arXiv:1705.02302*. [Online]. Available: <http://arxiv.org/abs/1705.02302>
- [13] Y. Levine, D. Yakira, N. Cohen, and A. Shashua, "Deep learning and quantum entanglement: Fundamental connections with implications to network design," in *Proc. 6th Int. Conf. Learn. Represent. (ICLR)*, Vancouver, BC, Canada, Apr./May 2018, pp. 1–46
- [14] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and J. I. Cirac, "Expressive power of tensor-network factorizations for probabilistic modeling," in *Proc. Annu. Conf. Neural Inf. Process. Syst. (NIPS)*, Vancouver, BC, Canada, Dec. 2019, pp. 1496–1508.
- [15] A. Anandkumar, R. Ge, D. J. Hsu, S. M. Kakade, and M. Telgarsky, "Tensor decompositions for learning latent variable models," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 2773–2832, 2014.
- [16] E. M. Stoudenmire and D. J. Schwab, "Supervised learning with tensor networks," in *Proc. Annu. Conf. Neural Inf. Process. Syst.*, Barcelona, Spain, Dec. 2016, pp. 4799–4807.
- [17] A. Novikov, M. Trofimov, and I. V. Oseledets, "Exponential machines," in *Proc. 5th Int. Conf. Learn. Represent. (ICLR)*, Toulon, France, Apr. 2017, pp. 1–13.
- [18] M. Blondel, M. Ishihata, A. Fujino, and N. Ueda, "Polynomial networks and factorization machines: New insights and efficient training algorithms," in *Proc. 33rd Int. Conf. Mach. Learn. (ICML)*, New York, NY, USA, Jun. 2016, pp. 850–858.
- [19] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [20] V. Dunjko and H. J. Briegel, "Machine learning & artificial intelligence in the quantum domain: A review of recent progress," *Rep. Prog. Phys.*, vol. 81, no. 7, Jul. 2018, Art. no. 074001.
- [21] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014, *arXiv:1411.4028*. [Online]. Available: <http://arxiv.org/abs/1411.4028>
- [22] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New J. Phys.*, vol. 18, no. 2, 2016, Art. no. 023023.
- [23] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," 2014, *arxiv:1802.0600*. [Online]. Available: <https://arxiv.org/abs/1802.0600>
- [24] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, Feb. 2019, Art. no. 040504.
- [25] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, "A generative modeling approach for benchmarking and training shallow quantum circuits," *npj Quantum Inf.*, vol. 5, no. 1, Dec. 2019.
- [26] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Phys. Rev. A, Gen. Phys.*, vol. 98, no. 3, Sep. 2018, Art. no. 032309.
- [27] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, "Towards quantum machine learning with tensor networks," *Quantum Sci. Technol.*, vol. 4, no. 2, 2019, Art. no. 024001.
- [28] E. Grant, M. Benedetti, S. Cao, A. Hallam, J. Lockhart, V. Stojevic, A. G. Green, and S. Severini, "Hierarchical quantum classifiers," *npj Quantum Inf.*, vol. 4, no. 1, pp. 1–8, Dec. 2018.
- [29] K. Fukushima, "Neural network model for a mechanism of pattern recognition unaffected by shift in position—Neocognitron," IECE, Jaipur, India, Tech. Rep. IECE J62-A(10), 1979, pp. 658–665.
- [30] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural Comput.*, vol. 1, no. 4, pp. 541–551, Dec. 1989.
- [31] P. Smolensky, "Information processing in dynamical systems: Foundations of harmony theory," in *Parallel Distributed Processing*, vol. 1. Cambridge, MA, USA: MIT Press, 1986.
- [32] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural Comput.*, vol. 14, no. 8, pp. 1771–1800, Aug. 2002.
- [33] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, "Simulation of quantum many-body systems with strings of operators and Monte Carlo tensor contractions," *Phys. Rev. Lett.*, vol. 100, no. 4, Jan. 2008, Art. no. 040501.
- [34] A. Sfondrini, J. Cerrillo, N. Schuch, and J. I. Cirac, "Simulating two- and three-dimensional frustrated quantum systems with string-bond states," *Phys. Rev. B*, vol. 81, no. 21, Jun. 2010, Art. no. 214426.
- [35] I. V. Oseledets, "Tensor-train decomposition," *SIAM J. Sci. Comput.*, vol. 33, no. 5, pp. 2295–2317, Jan. 2011.
- [36] A. Gendiar and T. Nishino, "Latent heat calculation of the three-dimensional $q = 3, 4$, and 5 potts models by the tensor product variational approach," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 65, Apr. 2002, Art. no. 046702.
- [37] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, "Ground-state properties of quantum many-body systems: Entangled-plaquette states and variational Monte Carlo," *New J. Phys.*, vol. 11, no. 8, 2009, Art. no. 083026.
- [38] H. J. Changlani, J. M. Kinder, C. J. Umrigar, and G. K.-L. Chan, "Approximating strongly correlated wave functions with correlator product states," *Phys. Rev. B*, vol. 80, no. 24, Dec. 2009, Art. no. 245116.
- [39] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, "Quantum entanglement in deep learning architectures," *Phys. Rev. Lett.*, vol. 122, no. 6, Feb. 2019, Art. no. 065301.
- [40] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, "Unsupervised generative modeling using matrix product states," *Phys. Rev. X*, vol. 8, no. 3, Jul. 2018, Art. no. 031012.
- [41] Y. Liu, X. Zhang, M. Lewenstein, and S.-J. Ran, "Entanglement-guided architectures of machine learning by quantum tensor network," 2018, *arXiv:1803.09111*. [Online]. Available: <http://arxiv.org/abs/1803.09111>
- [42] D. Liu, S.-J. Ran, P. Wittek, C. Peng, R. B. García, G. Su, and M. Lewenstein, "Machine learning by unitary tensor network of hierarchical tree structure," *New J. Phys.*, vol. 21, no. 7, 2019, Art. no. 073059.
- [43] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," 2017, *arXiv:1708.07747*. [Online]. Available: <http://arxiv.org/abs/1708.07747>
- [44] Y. LeCun, C. Cortes, and C. J. Burges. (1998). *MNIST Handwritten Digit Database*. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [45] K. J. Piczak, "Environmental sound classification with convolutional neural networks," in *Proc. IEEE 25th Int. Workshop Mach. Learn. Signal Process. (MLSP)*, Sep. 2015, pp. 1–6.
- [46] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (Adaptive Computation and Machine Learning). Cambridge, MA, USA: MIT Press, 2009.
- [47] B. J. Frey, "Extending factor graphs so as to unify directed and undirected graphical models," in *Proc. Conf. Uncertainty Artif. Intell.*, 2003, pp. 257–264.
- [48] Q. Zhao, M. Sugiyama, and A. Cichocki, "Learning efficient tensor representations with ring structure networks," 2017, *arXiv:1705.08286*. [Online]. Available: <http://arxiv.org/abs/1705.08286>

- [49] A. Critch and J. Morton, "Algebraic geometry of matrix product states," *SIGMA Symmetry Integrability Geometry Methods Appl.*, vol. 10, p. 95, 2014, doi: [10.3842/SIGMA.2014.095](https://doi.org/10.3842/SIGMA.2014.095).
- [50] M. Kliesch, D. Gross, and J. Eisert, "Matrix-product operators and states: NP-hardness and undecidability," *Phys. Rev. Lett.*, vol. 113, no. 16, Oct. 2014, Art. no. 160503.
- [51] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, "Equivalence of restricted Boltzmann machines and tensor network states," *Phys. Rev. B*, vol. 97, no. 8, Feb. 2018, Art. no. 085104.
- [52] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, "Neural-network quantum states, string-bond states, and chiral topological states," *Phys. Rev. X*, vol. 8, no. 1, Jan. 2018, Art. no. 011006.
- [53] E. Robeva and A. Seigal, "Duality of graphical models and tensor networks," 2017, *arXiv:1710.01437*. [Online]. Available: <http://arxiv.org/abs/1710.01437>
- [54] J. D. Biamonte, S. R. Clark, and D. Jaksch, "Categorical tensor network states," *AIP Adv.*, vol. 1, no. 4, Dec. 2011, Art. no. 042172.
- [55] S. R. Clark, "Unifying neural-network quantum states and correlator product states via tensor networks," 2017, *arXiv:1710.03545*. [Online]. Available: <http://arxiv.org/abs/1710.03545>
- [56] F. Verstraete and J. I. Cirac, "Renormalization algorithms for quantum-many body systems in two and higher dimensions," 2004, *arxiv:0407066*. [Online]. Available: <https://arxiv.org/abs/0407066>
- [57] D.-L. Deng, X. Li, and S. D. Sarma, "Quantum entanglement in neural network states," *Phys. Rev. X*, vol. 7, no. 2, May 2017, Art. no. 021021.
- [58] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for Boltzmann machines," *Cognit. Sci.*, vol. 9, no. 1, pp. 147–169, 1985.
- [59] H. Larochelle and Y. Bengio, "Classification using discriminative restricted Boltzmann machines," in *Proc. 25th Int. Conf. Mach. Learn. (ICML)*, Helsinki, Finland, Jun. 2008, pp. 536–543.
- [60] H. Larochelle, M. Mandel, R. Pascanu, and Y. Bengio, "Learning algorithms for the classification restricted Boltzmann machine," *J. Mach. Learn. Res.*, vol. 13, pp. 643–669, Mar. 2012.
- [61] S. Cheng, J. Chen, and L. Wang, "Information perspective to probabilistic modeling: Boltzmann machines versus born machines," *Entropy*, vol. 20, no. 8, p. 583, 2018.
- [62] J.-G. Liu and L. Wang, "Differentiable learning of quantum circuit born machines," *Phys. Rev. A, Gen. Phys.*, vol. 98, no. 6, Dec. 2018, Art. no. 062324.
- [63] J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," in *Proc. ACM Int. Conf. Multimedia (MM)*, Orlando, FL, USA, 2014, pp. 1041–1044.
- [64] V. Pestun and Y. Vlassopoulos, "Tensor network language model," 2017, *arXiv:1710.10248*. [Online]. Available: <http://arxiv.org/abs/1710.10248>
- [65] V. Pestun, J. Terilla, and Y. Vlassopoulos, "Language as a matrix product state," 2017, *arXiv:1711.01416*. [Online]. Available: <http://arxiv.org/abs/1711.01416>



IVAN GLASSER received the Master of Science degrees from Ecole Polytechnique, in 2012, and the University of Cambridge, in 2013, and the Ph.D. degree from the Technical University of Munich, in 2018. He is currently a Postdoctoral Researcher with the Max Planck Institute of Quantum Optics, Garching, Germany. His research interest includes simulation of complex quantum systems to the use of theoretical and numerical methods from physics in machine learning.



NICOLA PANCOTTI received the Master of Science degree from the University La Sapienza, Rome, in 2014. He is currently pursuing the Ph.D. degree with the Technical University of Munich and the Max Planck Institute of Quantum Optics, Garching, Germany. His research interests include dynamics of quantum systems, tensor network methods, and quantum machine learning.



J. IGNACIO CIRAC received the Ph.D. degree in theoretical physics from Complutense University, Madrid, in 1991. From 1991 to 1996, he was an Associate Professor with the University of Castilla-La Mancha and a Research Associate with the University of Colorado, Boulder. From 1996 to 2001, he was a Professor of theoretical physics with the University of Innsbruck. Since 2001, he has been a member of the Max Planck Society and the Director of the Max Planck Institute of Quantum Optics, Garching, Germany. In 2002, he became an Honorary Professor with the Technical University of Munich. He is one of the pioneers of the field of quantum computing and quantum information theory, and made major contributions in quantum optics and many-body physics. He is a member of the Spanish and German (Leopoldina) Academies of Sciences, and a Corresponding Member of the Austrian Academy of Sciences. He holds six honorary doctorships. He has received numerous awards, among them the Benjamin Franklin Medal, in 2010, the Wolf Prize in Physics, in 2013, and the Max Planck Medal of the German Physical Society, in 2018.

• • •