

# Learning Generative Models for Active Inference using Tensor Networks

Samuel T. Wauthier<sup>1</sup>, Bram Vanhecke<sup>2</sup>, Tim Verbelen<sup>1</sup>, and Bart Dhoedt<sup>1</sup>

<sup>1</sup> IDLab, Department of Information Technology at Ghent University – imec,  
Technologiepark-Zwijnaarde 126, B-9052 Ghent, Belgium  
`{firstname.lastname}@ugent.be`

<sup>2</sup> University of Vienna, Faculty of Physics and Faculty of Mathematics,  
Quantum Optics, Quantum Nanophysics and Quantum Information,  
Boltzmanngasse 5, 1090 Vienna, Austria  
`bram.andre.roland.vanhecke@univie.ac.at`

**Abstract.** Active inference provides a general framework for behavior and learning in autonomous agents. It states that an agent will attempt to minimize its variational free energy, defined in terms of beliefs over observations, internal states and policies. Traditionally, every aspect of a discrete active inference model must be specified by hand, i.e. by manually defining the hidden state space structure, as well as the required distributions such as likelihood and transition probabilities. Recently, efforts have been made to learn state space representations automatically from observations using deep neural networks. In this paper, we present a novel approach of learning state spaces using quantum physics-inspired tensor networks. The ability of tensor networks to represent the probabilistic nature of quantum states as well as to reduce large state spaces makes tensor networks a natural candidate for active inference. We show how tensor networks can be used as a generative model for sequential data. Furthermore, we show how one can obtain beliefs from such a generative model and how an active inference agent can use these to compute the expected free energy. Finally, we demonstrate our method on the classic T-maze environment.

**Keywords:** Active inference · Tensor networks · Generative modeling.

## 1 Introduction

Active inference is a theory of behavior and learning in autonomous agents [5]. An active inference agent selects actions based on beliefs about the environment in an attempt to minimize its variational free energy. As a result, the agent will try to reach its preferred state and minimize its uncertainty about the environment at the same time.

This scheme assumes that the agent possesses an internal model of the world and that it updates this model when new information, in the form of observations, becomes available. In current implementations, certain components of the model must be specified by hand. For example, the hidden space structure, as well

as transition dynamics and likelihood, must be manually defined. Deep active inference models deal with this problem by learning these parts of the model through neural networks [13,2].

Tensor networks, as opposed to neural networks, are networks constructed out of contractions between tensors. In recent years, tensor networks have found their place within the field of artificial intelligence. More specifically, Stoudenmire and Schwab [12] showed that it is possible to train these networks in an unsupervised manner to classify images from the MNIST handwritten digits dataset [8]. Importantly, tensor networks have shown to be valuable tools for generative modeling. Han et al. [6] and Cheng et al. [3] used tensor networks for generative modeling of the MNIST dataset, while Mencia Uranga and Lamacraft [9] used a tensor network to model raw audio.

Tensor networks, which were originally developed to represent quantum states in many-body quantum physics, are a natural candidate for generative models for two reasons. Firstly, they were developed in order to deal with the curse of dimensionality in quantum systems, where the dimensionality of the Hilbert space grows exponentially with the number of particles. Secondly, they are used to represent quantum states and are, therefore, inherently capable of representing uncertainty, or, in the case of active inference, beliefs. For example, contrary to neural networks, tensor networks do not require specifying a probability distribution for the hidden state variables or output variables. Furthermore, tensor networks can be exactly mapped to quantum circuits, which is important for the future of quantum machine learning [1].

In this paper, we present a novel approach to learning state spaces, likelihood and transition dynamics using the aforementioned tensor networks. We show that tensor networks are able to represent generative models of sequential data and how beliefs (i.e. probabilities) about observations naturally roll out of the model. Furthermore, we show how to compute the expected free energy for such a model using the sophisticated active inference scheme. We demonstrate this using the classic T-maze environment.

Section 2 elaborates on the inner workings of a tensor network and explains how to train such a model. Section 3 explains the environment and how we applied a tensor network for planning with active inference. In section 4, we present and discuss the results of our experiments on the T-maze environment. Finally, in section 5, we summarize our findings and examine future prospects.

## 2 Generative modeling with tensor networks

A generative model is a statistical model of the joint probability  $P(X)$  of a set of variables  $X = (X_1, X_2, \dots, X_n)$ . As previously mentioned, quantum states inherently contain uncertainty, i.e. they embody the probability distribution of a measurement of a system. It is natural, then, to represent a generative model as a quantum state. Quantum states can be mathematically described through a wave function  $\Psi(x)$  with  $x = (x_1, x_2, \dots, x_n)$  a set of real variables, such that

the probability of  $x$  is given by the Born rule:

$$P(X = x) = \frac{|\Psi(x)|^2}{Z}, \quad (1)$$

with  $Z = \sum_{\{x\}} |\Psi(x)|^2$ , where the summation runs over all possible realizations of the values of  $x$ .

Recent work [10] has pointed out that quantum states can be efficiently parameterized using tensor networks. The simplest form of tensor network is the matrix product state (MPS), also known as a tensor train [11]. When representing a quantum state as an MPS, the wave function can be parameterized as follows:

$$\Psi(x) = \sum_{\alpha_1} \sum_{\alpha_2} \sum_{\alpha_3} \dots \sum_{\alpha_{n-1}} A_{\alpha_1}^{(1)}(x_1) A_{\alpha_1 \alpha_2}^{(2)}(x_2) A_{\alpha_2 \alpha_3}^{(3)}(x_3) \dots A_{\alpha_{n-1}}^{(n)}(x_n). \quad (2)$$

Here, each  $A_{\alpha_{i-1} \alpha_i}^{(i)}(x_i)$  denotes a tensor of rank 2 (aside from the tensors on the extremities which are rank 1) which depends on the input variable  $x_i$ . This way, the wave function  $\Psi(x)$  is decomposed into a series of tensors  $A^{(i)}$ .

Importantly, each possible value of an input variable  $x_i$  must be associated with a vector of unit norm [12]. That is, each value that  $x_i$  can assume must be represented by a vector in a higher-dimensional feature space. Furthermore, to allow for a generative interpretation of the model, the feature vectors should be orthogonal [12]. This means that the vectors associated to the possible values of  $x_i$  will form an orthonormal basis of the aforementioned feature space. For a variable which can assume  $n$  discrete values, this feature space will be  $n$ -dimensional. The dimensionality of the space is referred to as the local dimension.

The unit norm and orthogonality conditions can be satisfied by defining a feature map  $\phi^{(i)}(x_i)$ , which maps each value onto a vector. For example, if  $x_i \in \{0, 1, 2\}$ , a possible feature map is the one-hot encoding of each value:  $(1, 0, 0)$  for 0,  $(0, 1, 0)$  for 1, and  $(0, 0, 1)$  for 2. The feature map  $\phi^{(i)}(x_i)$  allows us to rewrite the  $A^{(i)}(x_i)$  in Eq. 2 as

$$A_{\alpha_{i-1} \alpha_i}^{(i)}(x_i) = \sum_{\beta_i} T_{\alpha_{i-1} \beta_i \alpha_i}^{(i)} \phi_{\beta_i}^{(i)}(x_i), \quad (3)$$

where  $T_{\alpha_{i-1} \beta_i \alpha_i}^{(i)}$  is a tensor of rank 3. Here, we have further decomposed  $A^{(i)}$  into a contraction of tensor  $T^{(i)}$  and the feature vector  $\phi^{(i)}(x_i)$ . In graphical notation, the MPS (cf. Eq. 2) becomes:

$$\Psi(x) = \begin{array}{ccccccc} \textcircled{T^{(1)}} & \textcircled{T^{(2)}} & \textcircled{T^{(3)}} & \dots & \textcircled{T^{(n)}} \\ | & | & | & & | \\ \textcircled{\phi^{(1)}} & \textcircled{\phi^{(2)}} & \textcircled{\phi^{(3)}} & & \textcircled{\phi^{(n)}} \end{array}$$

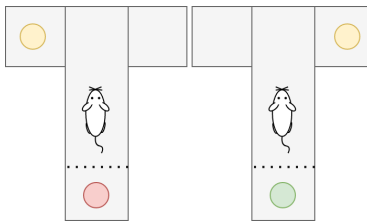
Given a data set, an MPS can be trained using a method based on the density matrix renormalization group (DMRG) algorithm [12]. This algorithm updates model parameters with respect to a given loss function by “sweeping” back and

forth across the MPS. In our case, we maximize the negative log-likelihood (NLL), i.e. we maximize the model evidence directly [6]. After training, the tensor network can be used to infer probability densities over unobserved variables by contracting the MPS with observed values. For a more in-depth discussion on tensor networks, we refer to the Appendix.

### 3 Environment

The environment used to test the generative model is the classic T-maze as presented by Friston et al. [5]. As the name suggests, the environment consists of a T-shaped maze and contains an artificial agent, e.g. a rat, and a reward, e.g. some cheese. The maze is divided into four locations: the center, the right branch, the left branch, and the cue location. The agent starts off in the center, while the reward is placed in either the left branch or the right branch, as depicted in Figure 1. Crucially, the agent does not know the location of the reward initially. Furthermore, once the agent chooses to go to either the left or the right branch, it is trapped and cannot leave. For the agent, the initial state of the world is uncertain. It does not know which type of world it is in: a world with the reward on the left, or a world with the reward on the right. In other words, it does not know its context. However, going to the cue reveals the location of the reward and enables the agent to observe the context. This resolves all ambiguity and allows the agent to make the correct decision.

The implementation of the environment was provided by pymdp [7]. In this package, the agent receives an observation with three modalities at every step: the location, the reward, and the context. The location can take on four possible values: center, right, left, and cue, and indicates which location the agent is currently in. The reward can take on three possible values: no reward, win, and loss. The “no reward” observation indicates that the agent received no information about the reward, while the “win” and “loss” observations indicate that the agent either received the reward or failed to obtain the reward, respectively. Logically, “no reward” can only be observed in the center and cue locations, while “win” and “loss” can only be observed in the left and right locations. Finally, the context



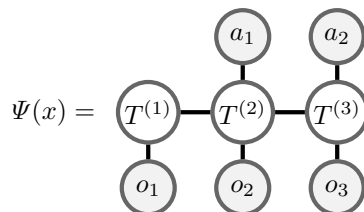
**Fig. 1.** Possible contexts of the T-maze environment as presented by Friston et al. [5]. The agent starts off in the center. The reward (yellow) is located in either the left branch (left image) or right branch (right image). The cue reveals where the reward is: red indicates the reward is on the left, green indicates the reward is on the right.

can take on two possible values: left and right. Whenever the agent is in locations “center”, “left” or “right”, the context observation will be randomly selected from “left” or “right” uniformly. Only when the agent is in the cue location, will the context observation yield the correct context. Further, the possible actions that the agent can take include: center, right, left, and cue, corresponding to which location the agent wants to go to.

We modified the above implementation slightly to better reflect the environment brought forth by Friston et al. [5]. In the original description, the agent is only allowed to perform two subsequent actions. Therefore, the number of time steps was limited to two. Furthermore, in the above implementation, the agent is able to leave the left and right branches of the T-maze. Thus, we prevented the agent from leaving whenever it chose to go to the left or right branches.

### 3.1 Modeling with tensor networks

The tensor network was adapted in order to accommodate the environment and be able to receive sequences of actions and observations as input. Firstly, the number of tensors was limited to the number of time steps. Secondly, each tensor received an extra index so that the network may receive both actions and observations. This led to the following network structure:



where we used  $a_i$  and  $o_i$  to denote the feature vectors corresponding to action  $a_i$  and observation  $o_i$ . Note that the first tensor did not receive an action input, since we defined that action  $a_i$  leads to observation  $o_{i+1}$ .

As mentioned in section 2, the feature maps  $\phi^{(i)}$  can generally be freely chosen as long as the resulting vectors are correctly normalized. However, it is useful to select feature maps which can easily be inverted, such that feature space vectors can readily be interpreted. In this case, since both observations and actions were discrete, one-hot encodings form a good option. The feature vectors for actions were one-hot encodings of the possible actions. The feature vectors for observations were one-hot encodings of the different combinations of observation modalities, i.e.  $4 \times 3 \times 2 = 24$  one-hot vectors corresponding to different combinations of the three different modalities.

In principle, there is nothing stopping us from learning feature maps, as long as the maps are correctly normalized. For practical purposes, the learning algorithm should make sure the feature maps are invertible. Whether feature maps should be learned before training the model or can be learned on-the-fly is an open question.

At this point, it is important to mention that the feature map is not chosen with the intent of imposing structure on the feature space based on prior knowledge of the observations (or actions). On the contrary, any feature map should assign a separate feature dimension to each possible observation, keeping the distance between that observation and all other observations within the feature space equal and maximal. To this end, the feature map can be thought of as being a part of the likelihood.

### 3.2 Planning with active inference

Once trained, the tensor network constructed in section 3.1 provides a generative model of the world. In theory, this model can be used for planning with active inference. At this point, it is important to remark that the network does not provide an accessible hidden state. While the bonds between tensors can be regarded as internal states, they are not normalized and, therefore, not usable. This poses a problem in the computation of the expected free energy, given by [5]

$$G(\pi) = \sum_{\tau} G(\pi, \tau) \quad (4)$$

$$G(\pi, \tau) = \mathbb{E}_{Q(o_{\tau}|s_{\tau}, \pi)}[\log Q(s_{\tau}|\pi) - \log P(s_{\tau}, o_{\tau}|\tilde{o}, \pi)] \quad (5)$$

$$\approx \underbrace{\text{D}_{\text{KL}}(Q(o_{\tau}|\pi) || P(o_{\tau}))}_{\text{expected cost}} + \underbrace{\mathbb{E}_{Q(s_{\tau}|\pi)}[\text{H}(Q(o_{\tau}|s_{\tau}, \pi))]}_{\text{expected ambiguity}}, \quad (6)$$

with hidden states  $s_{\tau}$ , observations  $o_{\tau}$  and policy  $\pi(\tau) = a_{\tau}$ , where the  $\sim$ -notation denotes a sequence of variables  $\tilde{x} = (x_1, x_2, \dots, x_{\tau-1})$  over time and  $P(o_{\tau})$  is the expected observation. This computation requires access to the hidden state  $s_{\tau}$  explicitly.

To remedy this, we suppose that the state  $s_{\tau}$  contains all the information gathered across actions and observations that occur at times  $t < \tau$ . Mathematically, we assume  $Q(s_{\tau}|\pi) \approx Q(o_{<\tau}|\pi)$  and  $Q(o_{\tau}|s_{\tau}, \pi) \approx Q(o_{\tau}|o_{<\tau}, \pi)$  with  $o_{<\tau} = (o_1, \dots, o_{\tau-1})$ . This way, we are able to approximate the expected ambiguity in Eq. 6. While these assumptions may give the impression that the calculation is computationally expensive, if the contraction with previous actions and observations has been performed once, it never has to be computed again, since the resulting tensor can be reused in subsequent calculations. At this point, the resulting tensor contains all the information from previous actions and observations.

When planning, we must re-evaluate the likelihood (and thus the expected free energy) for every time step, while imposing that the previous time steps are fixed. Indeed, we will perform sophisticated inference [4]. Under this scheme, the

expected free energy is given by

$$G(o_\tau, a_\tau) = \underbrace{\text{D}_{\text{KL}}(Q(o_{\tau+1}|a_{<\tau+1}) || P(o_{\tau+1})) + \mathbb{E}_{Q(s_{\tau+1}|a_{<\tau+1})}[\mathbb{H}(P(o_{\tau+1}|s_{\tau+1}))]}_{\text{expected free energy of next action}} + \underbrace{\mathbb{E}_{Q(a_{\tau+1}|o_{\tau+1})Q(o_{\tau+1}|a_{<\tau+1})}[G(o_{\tau+1}, a_{\tau+1})]}_{\text{expected free energy of subsequent actions}}, \quad (7)$$

$$Q(a_\tau|o_\tau) = \sigma[-G(o_\tau, a_\tau)] \quad (8)$$

This defines a tree search over actions and observations in the future.

## 4 Results and discussion

In this section, we demonstrate how the model’s beliefs shift over time. Later, we show how a tensor network agent behaves when planning under sophisticated inference.

The data set was constructed by including one of every possible path through the maze, i.e. 202 sequences of actions and observations. The model was trained over 500 epochs with a batch size of 10, where one epoch consisted of one right-to-left-to-right sweep per batch. The learning rate was set to  $10^{-4}$  and was further reduced by 10 % whenever the loss increased too much (i.e. by more than 0.5). Additionally, bonds started with 8 dimensions. The singular value cutoff point was set to 10 % of the largest singular value.

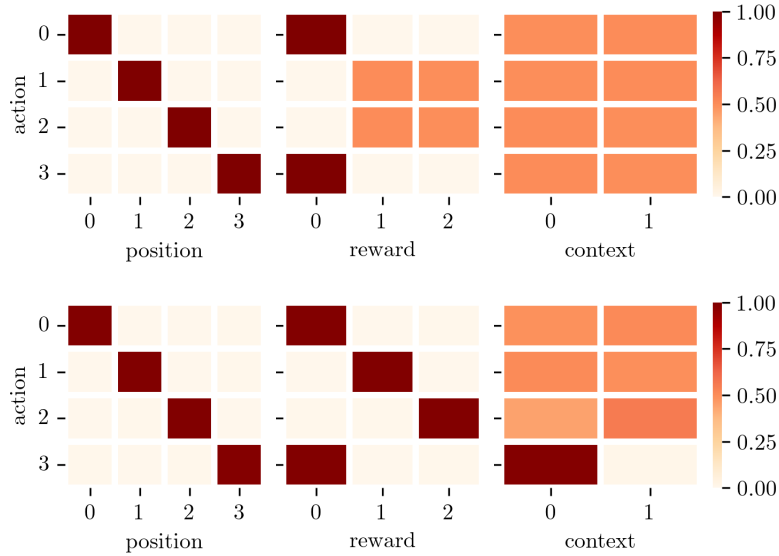
### 4.1 Belief shift

Since the initial observation  $o_1$  is always center position, no reward and context right or left with 50 % chance, we used the observation “center, no reward and context right” to obtain the beliefs in each case. The results are analogous in the case of “center, no reward and context left”.

Figure 2 (top) displays beliefs over  $o_2$  given  $a_1$ . From the results, it is clear that the agent does not know which reward it will receive, if it were to go to the left or right branch immediately. In addition, it does not know which context it will observe, even if the agent were to go towards the cue. Once the agent observes  $o_2$ , the beliefs shift. Figure 2 (bottom) shows beliefs over  $o_3$  given  $a_2$  when the agent has seen the cue with context “right“. Since the agent has seen the cue, it is very certain about the reward it will receive if it goes to the left or the right branch. If it stays in the cue location, it is also very certain that it will observe the same context again.

### 4.2 Action selection

With the outcome in section 4.1, we were able to perform action selection based on the sophisticated inference scheme described in section 3.2. For this, we used



**Fig. 2.** (top) Model beliefs over observation  $o_2$  given action  $a_1$  per modality. (bottom) Model beliefs for observation  $o_3$  given action  $a_2$  per modality, when the agent has observed the cue with context “right”. Actions 0, 1, 2 and 3 correspond to center, right, left and cue, respectively. Positions 0, 1, 2 and 3 correspond to center, right, left and cue location, respectively. Rewards 0, 1 and 2 correspond to no reward, win and loss, respectively. Context 0 and 1 correspond to right and left, respectively.

the following preferred observation per modality:

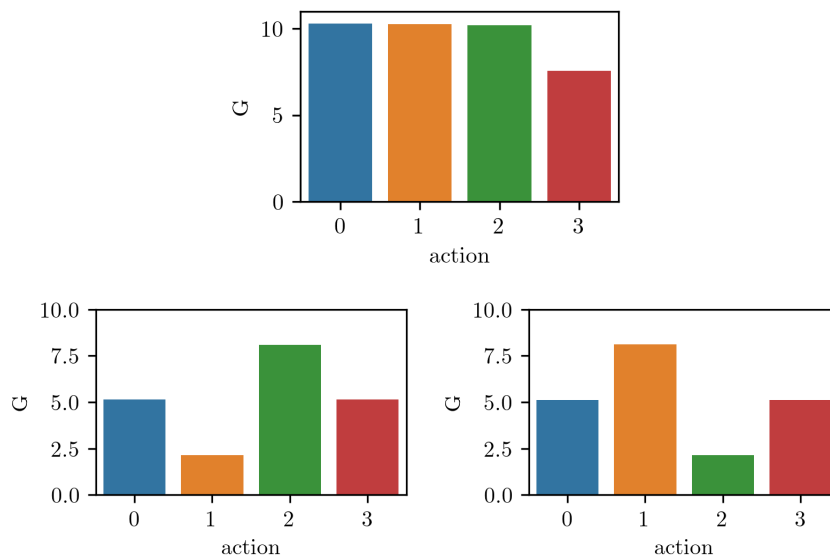
$$P(\text{position}) = \sigma([0 \ 0 \ 0 \ 0]), \quad P(\text{reward}) = \sigma([0 \ 3 \ -3]), \quad P(\text{context}) = \sigma([0 \ 0]). \quad (9)$$

Figure 3 (top) shows the expected free energy for action  $a_1$ . Given that the expected free energy is lowest for the action that brings the agent to the cue, the agent will choose to go to the cue in the first action. This is because, after observing the cue, the cue location provides a lower entropy on the context modality, as well as virtually 100% certainty on where the reward is located.

Figure 3 (bottom) shows the expected free energy for action  $a_2$  when the agent has chosen to go to the cue location and has observed either “cue, no reward, context right” or “cue, no reward, context left”. In this case, the agent will choose to go to either the left or the right branch, depending on the context it observed, i.e. context right will lead to action right and vice-versa.

The net result is that the agent will first go to the cue in order to resolve ambiguity and, subsequently, go to the branch with the reward.





**Fig. 3.** (top) Expected free energy for action  $a_1$ . (bottom) Expected free energy for action  $a_2$  when observation  $o_2$  was “cue, no reward, context right” and “cue, no reward, context left”, respectively.

## 5 Conclusion

We introduced a generative model based on tensor networks that is able to learn from sequential data. In addition, we showed how one can obtain beliefs from such a generative model and how a (sophisticated) active inference agent can use these to compute the expected free energy. Demonstration on the T-maze environment pointed out that such an agent is able to correctly select actions.

In the future, we plan to apply tensor networks to other environments, as well as make an in-depth comparison with neural networks, in order to better establish the benefits and drawbacks of the method. Moreover, we will adapt the network to allow sequences of random lengths and look into incorporating observations with continuous variables, which may also allow us to undo the assumptions made in section 3.2. Both these changes will broaden the range of applicability of generative models based on tensor networks.

## Acknowledgments

This research received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme. This work has received support from the European Union’s Horizon 2020 program through Grant No. 863476 (ERC-CoG SEQUAM).

## References




1. The tensor network, <https://tensornetwork.org/>, last accessed 21 June 2022
2. Çatal, O., Wauthier, S., De Boom, C., Verbelen, T., Dhoedt, B.: Learning generative state space models for active inference. *Frontiers in Computational Neuroscience* **14**, 103 (2020). <https://doi.org/10.3389/fncom.2020.574372>
3. Cheng, S., Wang, L., Xiang, T., Zhang, P.: Tree tensor networks for generative modeling. *Phys. Rev. B* **99**, 155131 (Apr 2019). <https://doi.org/10.1103/PhysRevB.99.155131>
4. Friston, K., Da Costa, L., Hafner, D., Hesp, C., Parr, T.: Sophisticated Inference. *Neural Computation* **33**(3), 713–763 (Mar 2021). [https://doi.org/10.1162/neco\\_a\\_01351](https://doi.org/10.1162/neco_a_01351)
5. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O’Doherty, J., Pezzulo, G.: Active inference and learning. *Neuroscience & Biobehavioral Reviews* **68**, 862–879 (2016). <https://doi.org/10.1016/j.neubiorev.2016.06.022>
6. Han, Z.Y., Wang, J., Fan, H., Wang, L., Zhang, P.: Unsupervised generative modeling using matrix product states. *Phys. Rev. X* **8**, 031012 (Jul 2018). <https://doi.org/10.1103/PhysRevX.8.031012>
7. Heins, C., Millidge, B., Demekas, D., Klein, B., Friston, K., Couzin, I.D., Tschantz, A.: pymdp: A python library for active inference in discrete state spaces. *Journal of Open Source Software* **7**(73), 4098 (2022). <https://doi.org/10.21105/joss.04098>
8. LeCun, Y., Cortes, C., Burges, C.J.C.: The mnist database of handwritten digits (1998), <http://yann.lecun.com/exdb/mnist/>, last accessed 10 June 2022
9. Mencia Uranga, B.n., Lamacraft, A.: Schrödinger: Generative modeling of raw audio as a continuously observed quantum state. In: Lu, J., Ward, R. (eds.) *Proceedings of the First Mathematical and Scientific Machine Learning Conference. Proceedings of Machine Learning Research*, vol. 107, pp. 74–106. PMLR (20–24 Jul 2020), <https://proceedings.mlr.press/v107/mencia-uranga20a.html>
10. Orús, R.: Tensor networks for complex quantum systems. *Nature Reviews Physics* **1**(9), 538–550 (Sep 2019). <https://doi.org/10.1038/s42254-019-0086-7>
11. Perez-Garcia, D., Verstraete, F., Wolf, M.M., Cirac, J.I.: Matrix product state representations. *Quantum Info. Comput.* **7**(5), 401–430 (Jul 2007)
12. Stoudenmire, E., Schwab, D.J.: Supervised learning with tensor networks. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016), <https://proceedings.neurips.cc/paper/2016/file/5314b9674c86e3f9d1ba25ef9bb32895-Paper.pdf>
13. Ueltzhöffer, K.: Deep active inference. *Biological Cybernetics* **112**(6), 547–573 (Dec 2018). <https://doi.org/10.1007/s00422-018-0785-7>

## Appendix 1 Notes on tensor networks

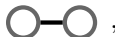
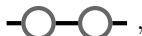

The summation over a common index as in Eq. 2 is also called a contraction. Performing the contraction between two tensors yields a new tensor with a rank equal to the sum of the ranks of the two contracted tensors minus two times the number of indices contracted over. That is, contracting two tensors of rank 2 over a single index gives a new tensor of rank 2, which is simply matrix multiplication:  $\sum_j A_{ij} B_{jk}$ . Similarly, contracting over the indices of a single tensor of rank 2 is simply the trace:  $\sum_i A_{ii}$ . In this sense, contraction is a generalization of these operations. Contracting indices in different ways gives rise to different structures. Examples of other possible networks are: tree tensor networks (TTN) and projected entangled pair states (PEPS).

Tensor networks can more easily be understood using their graphical notation. Each tensor is represented by a node, while contractions are represented by edges. Free edges, i.e. edges which do not connect two tensors, correspond to free indices which have not been summed. These can be used to represent sites in the network which are able to receive input or which can be used as input.

Some examples of tensors in graphical notation are:

- vector 
- matrix 
- rank-3 tensor 

Some examples of operations that can be represented by contractions are:

- dot product 
- matrix multiplication 
- trace 

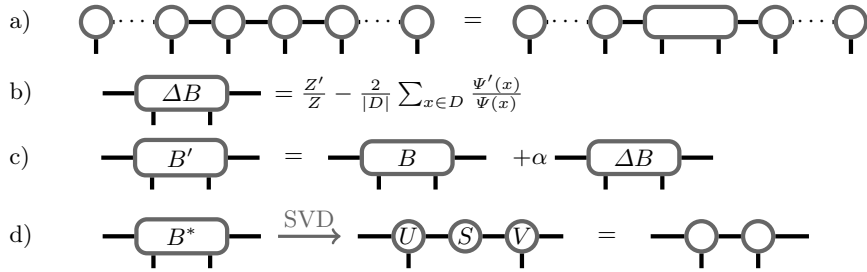
For a detailed account on tensor networks and their graphical notation, please refer to [1].

## Appendix 2 Training

The loss function must be chosen in such a way that the model captures the probability distribution of the data [6]. A straightforward method for estimating the parameters of a probability distribution is maximum likelihood estimation. In machine learning terms, this means we will optimize the parameters of the model with respect to the negative log-likelihood (NLL):

$$\mathcal{L} = \frac{1}{|D|} \sum_{x \in D} \log P(x), \quad (10)$$

where  $D$  denotes the data set. Through NLL minimization, the generative model becomes more similar to the probability distribution of the data.



**Fig. 4.** Training scheme for an MPS. a) Contraction of two adjacent tensors. b) Computing the update to the contracted tensor. c) Updating the contracted tensor. d) Decomposition of the contracted tensor using SVD.

Training proceeds as depicted in Figure 4. Firstly, tensor  $T^{(i)}$  and  $T^{(i+1)}$  are contracted to form the tensor  $B^{(i,i+1)}$ . The update to  $B^{(i,i+1)}$  is then computed using the loss function:

$$\Delta B^{(i,i+1)} = \frac{\partial \mathcal{L}}{\partial B_{\alpha_{i-1}\beta_i\alpha_{i+1}}^{(i,i+1)}} = \frac{Z'}{Z} - \frac{2}{|D|} \sum_{x \in D} \frac{\Psi'(x)}{\Psi(x)}, \quad (11)$$

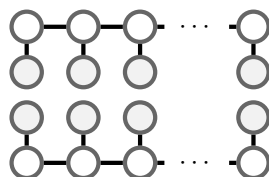
where  $Z' = 2 \sum_{x \in D} \Psi'(x)\Psi(x)$  and  $\Psi'$  is the derivative of  $\Psi$  with respect to  $B^{(i,i+1)}$ . Subsequently, the elements of  $B^{(i,i+1)}$  are adjusted by adding  $\Delta B^{(i,i+1)}$  multiplied by the learning rate. Finally, the newly computed  $B^{(i,i+1)}$  is decomposed into two tensors again. This decomposition is typically done through singular value decomposition (SVD), where the singular value matrix is then contracted with either the left or the right tensor, such that we are left with two tensors.

By starting this scheme at the leftmost tensor and iteratively moving one tensor to the right, the algorithm can update the entire MPS. Indeed, it is possible to update one tensor  $T_{\alpha_{i-1}\beta_i\alpha_i}^{(i)}$  at a time, however, the current method allows the dimensions of the indices  $\alpha_i$  (graphically, the edges connecting  $T^{(i)}$  nodes), the so-called bond dimensions, to vary during training. This is made possible by truncated SVD, which truncates dimensions with singular values that fall beneath some manually specified threshold. Truncating dimensions with small singular values can be interpreted as truncating less informative dimensions. As a result, truncated SVD ensures that the model remains as small as possible, while containing the most information. Moreover, the size of the model will vary depending on how much information it must learn.

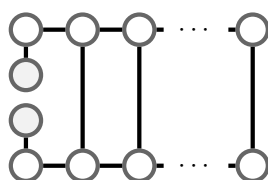
### Appendix 3 Computing probabilities with tensor networks

One of the benefits of tensor networks is that we can easily obtain exact joint (and conditional) probability distributions without requiring parameterization.

After training the model, given that the network is correctly normalized ( $Z = 1$ ), the joint probability distribution is given by

$$P(x_1, x_2, x_3, \dots, x_n) =$$


where we have omitted tensor labels for simplicity. Marginal distributions can be found by discarding the offending variables:

$$P(x_1) = \sum_{\{x\} \setminus \{x_1\}} P(x_1, x_2, x_3, \dots, x_n) =$$


where the sum over the variables  $x_2, x_3, \dots, x_n$  is equivalent to contracting the matching tensors. Finally, conditional probability distributions can be found by combining the previous results:

$$P(x_2, x_3, \dots, x_n | x_1) = \frac{P(x_1, x_2, x_3, \dots, x_n)}{P(x_1)}. \quad (12)$$

## Appendix 4 Physical intuition

In order to garner a feeling for the physics and mathematics used throughout this work, this section describes in (mostly) words what the physical meaning of the constituents of the tensor network is.

Let  $x$  be an observable, i.e. a quantity that can be measured (or observed). Examples of such physical quantities are position and momentum. Furthermore, let  $\{0, 1, 2\}$  be the set of values that  $x$  can assume.

In quantum mechanics, the set of values that an observable can assume are eigenvalues. This entails that there is a set of eigenvectors corresponding to these eigenvalues. In turn, the eigenvectors form an eigenbasis of the state space. In other words, every value that  $x$  can assume has a corresponding vector and each of these vectors is a basis vector of the state space, meaning that each different value is represented in the state space by a different dimension. For example, we may have the vectors  $(1, 0, 0)$  for 0,  $(0, 1, 0)$  for 1, and  $(0, 0, 1)$  for 2.

Further, an MPS is typically used to represent a state within the state space. This means that the MPS represents a (multi)vector. This vector is not necessarily one of the eigenvectors mentioned earlier, but it can be virtually any vector within the state space. When learning the parameters of an MPS, it is rotated and stretched in such a way that it represents the data.

If an MPS does not necessarily coincide with any of the eigenvectors, what value will it produce when a measurement is performed? It will produce any of the eigenvalues with a certain probability. These probabilities are given by the square of the inner product between the MPS and each eigenvector. In our example, say the MPS was represented by the vector  $(0.55, 0.55, 0.63)$ , we would measure 0 with 30% probability, 1 with 30% probability and 2 with 40% probability.