

Neural Decoder for Topological Codes

Giacomo Torlai and Roger G. Melko

*Department of Physics and Astronomy, University of Waterloo, Ontario N2L 3G1, Canada
and Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

(Received 20 October 2016; published 18 July 2017)

We present an algorithm for error correction in topological codes that exploits modern machine learning techniques. Our decoder is constructed from a stochastic neural network called a Boltzmann machine, of the type extensively used in deep learning. We provide a general prescription for the training of the network and a decoding strategy that is applicable to a wide variety of stabilizer codes with very little specialization. We demonstrate the neural decoder numerically on the well-known two-dimensional toric code with phase-flip errors.

DOI: 10.1103/PhysRevLett.119.030501

Introduction.—Much of the success of modern machine learning stems from the flexibility of a given neural network architecture to be employed for a multitude of different tasks. This generalizability means that neural networks can have the ability to infer structure from vastly different data sets with only a change in optimal hyperparameters. For this purpose, the machine learning community has developed a set of standard tools, such as fully connected feed forward networks [1] and Boltzmann machines [2]. Specializations of these underlie many of the more advanced algorithms, including convolutional networks [3] and deep learning [4,5], encountered in real-world applications such as image or speech recognition [6].

These machine learning techniques may be harnessed for a multitude of complex tasks in science and engineering [7–17]. An important application lies in quantum computing. For a quantum logic operation to succeed, noise sources that lead to decoherence in a qubit must be mitigated. This can be done through some type of quantum error correction—a process where the logical state of a qubit is encoded redundantly so that errors can be corrected before they corrupt it [18]. A leading candidate for this is the implementation of fault-tolerant hardware through surface codes, where a logical qubit is stored as a topological state of an array of physical qubits [19]. Random errors in the states of the physical qubits can be corrected before they proliferate and destroy the logical state. The quantum error correction protocols that perform this correction are termed “decoders,” and must be implemented by classical algorithms running on conventional computers [20,21].

In this Letter we demonstrate how one of the simplest stochastic neural networks for unsupervised learning, the restricted Boltzmann machine [22], can be used to construct a general error-correction protocol for stabilizer codes. Given a syndrome, defined by a measurement of the end points of an (unknown) chain of physical qubit errors, we use our Boltzmann machine to devise a protocol with the goal of correcting errors without corrupting the

logical bit. Our decoder works for generic degenerate stabilizer codes that have a probabilistic relation between the syndrome and errors, which does not have to be *a priori* known. Importantly, it is very simple to implement, requiring no specialization regarding code locality, dimension, or structure. We test our decoder numerically on a simple two-dimensional surface code with phase-flip errors.

The 2D toric code.—Most topological codes can be described in terms of the stabilizer formalism [23]. A stabilizer code is a particular class of error-correcting code characterized by a protected subspace \mathcal{C} defined by a stabilizer group \mathcal{S} . The simplest example is the 2D toric code, first introduced by Kitaev [24]. Here, the quantum information is encoded into the homological degrees of freedom, with topological invariance given by the first homology group [25]. The code features N qubits placed on the links of an $L \times L$ square lattice embedded on a torus. The stabilizer group is $\mathcal{S} = \{\hat{Z}_p, \hat{X}_v\}$, where the plaquette and vertex stabilizers are defined respectively as $\hat{Z}_p = \bigotimes_{\ell \in p} \hat{\sigma}_\ell^z$ and $\hat{X}_v = \bigotimes_{\ell \in v} \hat{\sigma}_\ell^x$ with $\hat{\sigma}_\ell^z$ and $\hat{\sigma}_\ell^x$ acting, respectively, on the links contained in the plaquette p and the links connected to the vertex v . There are two encoded logical qubits, manipulated by logical operators $\hat{Z}_L^{(1,2)}$ as $\hat{\sigma}^z$ acting on the noncontractible loops on the real lattice and logical $\hat{X}_L^{(1,2)}$ as the noncontractible loops on the dual lattice (Fig. 1).

Given a reference state $|\psi_0\rangle \in \mathcal{C}$, let us consider the simple phase-flip channel described by a Pauli operator where $\hat{\sigma}^z$ is applied to each qubit with probability p_{err} . This operator can be efficiently described by a mapping between the links and \mathbb{Z}_2 , called an error chain \mathbf{e} , whose boundary is called a syndrome $\mathcal{S}(\mathbf{e})$. In an experimental implementation, only the syndrome (and not the error chain) can be measured. Error correction (decoding) consists of applying a recovery operator whose chain \mathbf{r} generates the same syndrome, $\mathcal{S}(\mathbf{e}) = \mathcal{S}(\mathbf{r})$. The recovery succeeds only if the combined operation is described by a cycle (i.e., a chain




FIG. 1. Several operations on a 2D toric code. Logical operators $\hat{Z}_L^{(1)}$ and $\hat{Z}_L^{(2)}$ (orange) are nontrivial cycles on the real lattice. A physical error chain e (purple) and its syndrome $S(e)$ (black squares). A recovery chain r' (green), with the combined operator on the cycle $e \oplus r'$ being a product of stabilizers $\hat{Z}_\alpha \hat{Z}_\beta \hat{Z}_\gamma$ (recovery success). A recovery chain r'' (red) whose cycle has nontrivial homology and acts on the code state as $\hat{Z}_L^{(1)}$ (logical failure).

with no boundaries) $e \oplus r$ that belongs to the trivial homology class h_0 , describing contractible loops on the torus. On the other hand, if the cycle belongs to a nontrivial homology class (being noncontractible on the torus), the recovery operation directly manipulates the encoded logical information, leading to a logical failure (Fig. 1).

Several decoders have been proposed for the 2D toric code, based on different strategies [26–30]. Maximum likelihood decoding consists of finding a recovery chain r with the most likely homology class [31,32]. A different recovery strategy, designed to reduce computational complexity, consists of generating the recovery chain r compatible with the syndrome simply by using the minimum number of errors. Such a procedure, called minimum weight perfect matching [33] (MWPM), has the advantage that it can be performed without the knowledge of the error probability p_{err} . This algorithm is however suboptimal (with lower threshold probability [25]) since it does not take into account the high degeneracy of the error chains given a syndrome.

The neural decoder.—Neural networks are commonly used to extract features from raw data in terms of probability distributions. In order to exploit this for error correction, we first build a data set made of error chains and their syndromes $\mathcal{D} = \{e, S\}$, and train a neural network to model the underlying probability distribution $p_{\text{data}}(e, S)$. Our goal is to then generate error chains to use for the recovery. We use a generative model called a Boltzmann machine, a powerful stochastic neural network widely used in the pretraining of the layers of deep neural networks [34,35]. The network architecture features three layers of stochastic binary neurons, the syndrome layer $S \in \{0, 1\}^{N/2}$, the error layer $e \in \{0, 1\}^N$, and one hidden layer $h \in \{0, 1\}^{n_h}$ (Fig. 2). Symmetric edges connect both




FIG. 2. The neural decoder architecture. The hidden layer h is fully connected to the syndrome and error layers S and e with weights U and W , respectively.

the syndrome and the error layer with the hidden layer. We point out that this network is equivalent to a traditional bilayer restricted Boltzmann machine, where we have here divided the visible layer into two separate layers for clarity. The weights on the edges connecting the network layers are given by the matrices U and W . Moreover, we also add external fields b , c , and d coupled to the every neuron in each layer. The probability distribution that the probabilistic model associates with this graph structure is the Boltzmann distribution [36]

$$p_\lambda(e, S, h) = \frac{1}{Z_\lambda} e^{-E_\lambda(e, S, h)}, \quad (1)$$

where $Z_\lambda = \text{Tr}_{\{h, S, e\}} e^{-E_\lambda(e, S, h)}$ is the partition function, $\lambda = \{U, W, b, c, d\}$ is the set of parameters of the model, and the energy is

$$E_\lambda(e, S, h) = -\sum_{ik} U_{ik} h_i S_k - \sum_{ij} W_{ij} h_i e_j - \sum_j b_j e_j - \sum_i c_i h_i - \sum_k d_k S_k. \quad (2)$$

The joint probability distribution over (e, S) is obtained after integrating out the hidden variables from the full distribution,

$$p_\lambda(e, S) = \sum_h p_\lambda(e, S, h) = \frac{1}{Z_\lambda} e^{-\mathcal{E}_\lambda(e, S)}, \quad (3)$$

where the effective energy $\mathcal{E}_\lambda(e, S)$ can be computed exactly. Moreover, given the structure of the network, the conditional probabilities $p_\lambda(e|h)$, $p_\lambda(S|h)$, and $p_\lambda(h|e, S)$ are also known exactly. The training of the machine consists of tuning the parameters λ until the model probability $p_\lambda(e, S)$ becomes close to the target distribution $p_{\text{data}}(e, S)$ of the data set. This translates into solving an optimization problem over the parameters λ by minimizing the distance between the two distributions, defined as the Kullback-Leibler divergence, $\mathbb{KL} \propto -\sum_{(e, S) \in \mathcal{D}} \log p_\lambda(e, S)$. Details about the Boltzmann

machine and its training algorithm are reported in the Supplemental Material [37].

We now discuss the decoding algorithm, which proceeds assuming that we successfully learned the distribution $p_\lambda(\mathbf{e}, \mathbf{S})$. Given an error chain \mathbf{e}_0 with syndrome \mathbf{S}_0 we wish to use the Boltzmann machine to generate an error chain compatible with \mathbf{S}_0 to use for the recovery. To achieve this goal we separately train networks on different data sets obtained from different error regimes p_{err} . Assuming we know the error regime that generated \mathbf{e}_0 , the recovery procedure consists of sampling a recovery chain from the distribution $p_\lambda(\mathbf{e}|\mathbf{S}_0)$ given by the network trained at the same probability p_{err} of \mathbf{e}_0 . Although the Boltzmann machine does not learn this distribution directly, by sampling the error and hidden layers while keeping the syndrome layer fixed to \mathbf{S}_0 , since $p_\lambda(\mathbf{e}, \mathbf{S}_0) = p_\lambda(\mathbf{e}|\mathbf{S}_0)p(\mathbf{S}_0)$, we are enforcing sampling from the desired conditional distribution. An advantage of this procedure over decoders that employ conventional Monte Carlo calculations [28,29] on specific stabilizer codes is that specialized sampling algorithms tied to the stabilizer structure, or multicanonical methods such as parallel tempering, are not required. Finally, note that the assumption of perfect learning is not critical, since the above sampling routine can be modified with an extra rejection step as discussed in Ref. [14] to ensure sampling occurs from the proper physical distribution.

An error correction procedure can be defined as follows (Algorithm 1): we first initialize the machine into a random state of the error and hidden layers (see Fig. 2) and to \mathbf{S}_0 for the syndrome layer. We then let the machine equilibrate by repeatedly performing block Gibbs sampling. After some amount of equilibration steps, we begin checking the syndrome of the error state \mathbf{e} in the machine and, as soon as $\mathbf{S}(\mathbf{e}) = \mathbf{S}_0$, we select it for the recovery operation. If such a condition is not met before a fixed amount of sampling steps, the recovery attempt is stopped and considered failed. This condition makes the precise computational requirements of the algorithm ill defined, since the cutoff time can always be increased resulting in better performance for a higher computational cost.

Algorithm 1 Neural Decoding Strategy

- | | |
|--|-------------------------|
| 1: \mathbf{e}_0 : physical error chain | |
| 2: $\mathbf{S}_0 = \mathbf{S}(\mathbf{e}_0)$ | ▷Syndrome Extraction |
| 3: Boltzmann machine = $\{\mathbf{e}, \mathbf{S} = \mathbf{S}_0, \mathbf{h}\}$ | ▷Network Initialization |
| 4: <i>while</i> $\mathbf{S}(\mathbf{e}) \neq \mathbf{S}_0$ <i>do</i> | ▷Sampling |
| 5: Sample $\mathbf{h} \sim p(\mathbf{h} \mathbf{e}, \mathbf{S}_0)$ | |
| 6: Sample $\mathbf{e} \sim p(\mathbf{e} \mathbf{h})$ | |
| 7: <i>end while</i> | |
| 8: $\mathbf{r} = \mathbf{e}$ | ▷Decoding |
-

Results.—We train neural networks in different error regimes by building several data sets $\mathcal{D}_p = \{\mathbf{e}_k, \mathbf{S}_k\}_{k=1}^M$ at elementary error probabilities $p = \{0.5, 0.6, \dots, 0.15\}$ of

the phase-flip channel. For a given error probability, the network hyperparameters are individually optimized via a grid search (for details see the Supplemental Material [37]). Once training is complete, we perform decoding following the procedure laid out in Algorithm 1. We generate a test set $\mathcal{T}_p = \{\mathbf{e}_k\}_{k=1}^M$ and for each error chain $\mathbf{e}_k \in \mathcal{T}_p$, after a suitable equilibration time (usually $N_{\text{eq}} \propto 10^2$ sampling steps), we collect the first error chain \mathbf{e} compatible with the original syndrome, $\mathbf{S}(\mathbf{e}) = \mathbf{S}(\mathbf{e}_k)$. We use this error chain for the recovery, $\mathbf{r}^{(k)} = \mathbf{e}$. Importantly, error recovery with $\mathbf{r}^{(k)}$ chosen from the first compatible chain means that the cycle $\mathbf{e}_k + \mathbf{r}^{(k)}$ is sampled from a distribution that includes all homology classes. By computing the Wilson loops on the cycles we can measure their homology class. This allows us to gauge the accuracy of the decoder in term of the logical failure probability, defined as $P_{\text{fail}} = n_{\text{fail}}/M$, where n_{fail} is the number of cycles with nontrivial homology. Because of the fully connected architecture of the network, and the large complexity of the probability distribution arising from the high degeneracy of error chains given a syndrome, we found that the data set size required to accurately capture the underlying statistics must be relatively large ($|\mathcal{D}_p| \propto 10^5$). In Fig. 3 we plot the logical failure probability P_{fail} as a function of the elementary error probability for the neural decoding scheme. We note that at low p_{err} , our logical failure probabilities follow the expected [38] scaling form $p_{\text{err}}^{L/2}$ (not plotted).

To compare our numerical results we also perform error correction using the recovery scheme given by MWPM [39]. This algorithm creates a graph whose vertices corresponds to the syndrome and the edges connect each vertex with a weight equal to the Manhattan distance (the number of links connecting the vertices in the original square lattice). MWPM then finds an optimal matching of all the vertices pairwise using the minimum weight, which corresponds to the minimum number of edges in the lattice [40]. Figure 3 displays the comparison between a MWPM

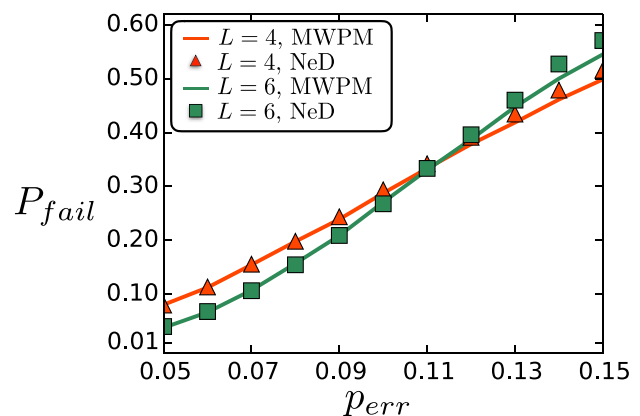


FIG. 3. Logical failure probability as a function of elementary error probability for MWPM (lines) and the neural decoder (markers) of size $L = 4$ (red) and $L = 6$ (green).

decoder (line) and our neural decoder (markers). As is evident, the neural decoder has an almost identical logical failure rate for error probabilities below the threshold ($p_{\text{err}} \approx 10.9$ [25]), yet a significant higher probability above. Note that by training the Boltzmann machine on different data sets we have enforced in the neural decoder a dependence on the error probability. This is in contrast to MWPM, which is performed without such knowledge. Another key difference is that the distributions learned by the Boltzmann machine contain the entropic contribution from the high degeneracy of error chains, which is directly encoded into the data sets. It will be instructive to explore this further, to determine whether the differences in Fig. 3 come from inefficiencies in the training, the different decoding model of the neural network, or both. Finite-size scaling on larger L will allow calculation of the threshold defined by the neural decoder.

In the above algorithm, which amounts to a simple and practical implementation of the neural decoder, our choice to use the first compatible chain for error correction means that the resulting logical operation is sampled from a distribution that includes all homology classes. This is illustrated in Fig. 4, where we plot the histogram of the homology classes for several different elementary error probabilities. Accordingly, our neural decoder can easily be modified to perform maximum likelihood (ML) optimal decoding. For a given syndrome, instead of obtaining only one error chain to use in decoding, one could sample many error chains and build up the histogram of homology classes with respect to any reference error state. Then, choosing the recovery chain from the largest histogram bin will implement, by definition, ML decoding. Although the computational cost of this procedure will clearly be expensive using the current fully connected restricted Boltzmann machine, it would be interesting to explore




FIG. 4. Histogram of the homology classes returned by our neural decoder for various elementary error probabilities p_{err} . The green bars represent the trivial homology class h_0 corresponding to contractible loops on the torus. The other three classes correspond, respectively, to the logical operations $\hat{Z}_L^{(1)}$, $\hat{Z}_L^{(2)}$, and $\hat{Z}_L^{(1)}\hat{Z}_L^{(2)}$.

specializations of the neural network architecture in the future to see how its performance may compare to other ML decoding algorithms [31].

Conclusions.—We have presented a decoder for topological codes using a simple algorithm implemented with a restricted Boltzmann machine, a common neural network used in many machine learning applications. Our neural decoder is easy to program using standard machine learning software libraries and training techniques, and relies on the efficient sampling of error chains distributed over all homology classes. Numerical results show that our decoder has a logical failure probability that is close to MWPM, but not identical, a consequence of our neural network being trained separately at different elementary error probabilities. This leads to the natural question of the relationship between the neural decoder and optimal decoding, which could be explored further by a variation of our algorithm that implements maximum likelihood decoding.

In its current implementation, the Boltzmann machine is restricted within a given layer of neurons, but fully connected between layers. This means that our decoder does not depend on the specific geometry used to implement the code, or on the structure of the stabilizer group; it is trained simply using a raw data input vector, with no information on locality or dimension. Such a high degree of generalizability, which is one of the core advantages of this decoder, also represents a challenge for investigating bigger systems. For example, a bottleneck in our scheme to decode larger sizes is finding an error chain compatible with the syndrome within a reasonable cutoff time.

In order to scale up our system sizes on the 2D toric code (as required, e.g., to calculate the threshold), one could relax some of the general fully connected structure of the network, and specialize it to accommodate the specific details of the code. Geometric specialization such as this has been explicitly demonstrated to improve the representational efficiency of neural networks in the case of the toric code [8,13]. This specialization should be explored in detail, before comparison of computational efficiency are made between our neural decoder, MWPM, and other decoding schemes. Note that, even with moderate specialization, the neural decoder as we have presented it above can immediately be extended to other choices of error models [41], such as the more realistic case of imperfect syndrome measurement [42], or transferred to other topological stabilizer codes, such as color codes [43,44]. We also point out that the training of the networks is performed off-line and has to be carried out only once. As such, the high computational cost of the training need not be considered when evaluating the decoder computational efficiency for any of these examples.

Finally, it would be interesting to explore the improvements in performance obtained by implementing standard tricks in machine learning, such as convolutions, adaptive optimization algorithms, or the stacking of multiple

Boltzmann machines into a network with deep structure. Given the rapid advancement of machine learning technology within the world's information industry, we expect that such tools will be the obvious choice for the real-world implementation of decoding schemes on future topologically fault-tolerant qubit hardware.

The authors thank J. Carrasquilla, D. Gottesman, M. Hastings, C. Herdmann, B. Kulchytskyy, M. Mariantoni, and D. Poulin for enlightening discussions. This research was supported by NSERC, the CRC program, the Ontario Trillium Foundation, Perimeter Institute for Theoretical Physics, and the National Science Foundation under Grant No. NSF PHY-1125915. Simulations were performed on resources provided by SHARCNET. Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation.

-
- [1] K. Hornik, M. Stinchcombe, and H. White, *Neural Netw.* **2**, 359 (1989).
- [2] R. Salakhutdinov, Technical Report UTML, Department of Computer Science, University of Toronto, 2008, p. 002.
- [3] A. Krizhevsky, I. Sutskever, and G. Hinton, *Proceeding Advances in Neural Information Processing Systems (NIP2012)*, Lake Tahoe (2012), Vol. **25**, p. 1090.
- [4] G. Hinton, *Trends Cognit. Sci.* **11**, 428 (2007).
- [5] Y. LeCun, Y. Bengio, and G. Hinton, *Nature (London)* **521**, 436 (2008).
- [6] G. Hinton *et al.*, *IEEE Signal Process. Mag.* **29**, 82 (2012).
- [7] G. Torlai and R. G. Melko, *Phys. Rev. B* **94**, 165134 (2016).
- [8] J. Carrasquilla and R. G. Melko, *Nat. Phys.* **13**, 431 (2017).
- [9] L. Wang, *Phys. Rev. B* **94**, 195105 (2016).
- [10] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [11] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, [arXiv:1608.07848](https://arxiv.org/abs/1608.07848).
- [12] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, [arXiv:1609.02552](https://arxiv.org/abs/1609.02552).
- [13] D.-L. Deng, X. Li, and S. D. Sarma, [arXiv:1609.09060](https://arxiv.org/abs/1609.09060).
- [14] L. Huang and L. Wang, *Phys. Rev. B* **95**, 035105 (2017).
- [15] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, *Phys. Rev. B* **95**, 041101 (2017).
- [16] E. M. Stoudenmire and D. J. Schwab, [arXiv:1605.05775](https://arxiv.org/abs/1605.05775).
- [17] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. G. Melko, and G. Carleo, [arXiv:1703.05334](https://arxiv.org/abs/1703.05334).
- [18] D. Nigg, M. Mueller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt, *Science* **345**, 302 (2014).
- [19] H. Bombin, in *Quantum Error Correction*, edited by D. A. Lidar and T. A. Brun (Cambridge University Press, Cambridge, England, 2013), Chap. 19.
- [20] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Phys. Rev. A* **86**, 032324 (2012).
- [21] H. Bombin and M. A. Martin-Delgado, *J. Phys. A* **42**, 095302 (2009).
- [22] G. Hinton, *Neural Networks: Tricks of the Trade* (Springer, Berlin Heidelberg, 2012), p. 599.
- [23] D. Gottesman, [arXiv:quant-ph/9705052](https://arxiv.org/abs/quant-ph/9705052).
- [24] A. Y. Kitaev, *Ann. Phys. (Amsterdam)* **303**, 2 (2003).
- [25] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *J. Math. Phys. (N.Y.)* **43**, 4452 (2002).
- [26] G. Duclos-Cianci and D. Poulin, *Phys. Rev. Lett.* **104**, 050504 (2010).
- [27] G. Duclos-Cianci and D. Poulin, *Quant. Inf. Comp.* **14**, 0721 (2014).
- [28] J. R. Wootton and D. Loss, *Phys. Rev. Lett.* **109**, 160503 (2012).
- [29] A. Hutter, J. R. Wootton, and D. Loss, *Phys. Rev. A* **89**, 022326 (2014).
- [30] A. Fowler, [arXiv:1310.0863](https://arxiv.org/abs/1310.0863).
- [31] S. Bravyi, M. Suchara, and A. Vargo, *Phys. Rev. A* **90**, 032326 (2014).
- [32] B. Heim, K. M. Svore, and M. B. Hastings, [arXiv:1609.06373](https://arxiv.org/abs/1609.06373).
- [33] J. Edmonds, *Can. J. Math.* **17**, 449 (1965).
- [34] G. Hinton, S. Osindero, and Y. Teh, *Neural Comput.* **18**, 1527 (2006).
- [35] R. Salakhutdinov and I. Murray, in *ICML-08 Proceedings of the 25th international conference on Machine learning* (ACM, New York, 2008), p. 872.
- [36] A. Fischer and C. Igel, *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Springer, Berlin Heidelberg, 2012), p. 14.
- [37] Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.119.030501> for details about the Boltzmann machine and its training.
- [38] F. H. E. Watson and S. D. Barrett, *New J. Phys.* **16**, 093045 (2014).
- [39] V. Kolmogorov, *Math. Prog. Comp.* **1**, 43 (2002).
- [40] A. G. Fowler, A. C. Whiteside, and L. C. L. Hollenberg, *Phys. Rev. Lett.* **108**, 180501 (2012).
- [41] E. Novais and E. R. Mucciolo, *Phys. Rev. Lett.* **110**, 010502 (2013).
- [42] C. Wang, J. Harrington, and J. Preskill, *Ann. Phys. (Amsterdam)* **303**, 31 (2003).
- [43] H. G. Katzgraber, H. Bombin, and M. A. Martin-Delgado, *Phys. Rev. Lett.* **103**, 090501 (2009).
- [44] B. J. Brown, N. H. Nickerson, and D. E. Browne, *Nat. Commun.* **7**, 12302 (2016).