

## Determination of the Lowest-Energy States for the Model Distribution of Trained Restricted Boltzmann Machines Using a 1000 Qubit D-Wave 2X Quantum Computer

**Yaroslav Koshka**

*ykoshka@ece.msstate.edu*

*Department of Electrical and Computer Engineering, Emerging Materials Research Laboratory, HPC<sup>2</sup> Distributed Analytics and Security Institute, Mississippi State University, Mississippi State, MS 39762, U.S.A.*

**Dilina Perera**

*dilina@dasi.msstate.edu*

**Spencer Hall**

*jsh278@msstate.edu*

*HPC<sup>2</sup> Distributed Analytics and Security Institute, Mississippi State University, Mississippi State, MS 39762, U.S.A.*

**M. A. Novotny**

*man40@msstate.edu*

*Department of Physics and Astronomy, HPC<sup>2</sup> Center for Computational Sciences, Mississippi State University, Mississippi State, MS 39762, U.S.A., and Faculty of Mathematica and Physics, Charles University Ke Karlovu 5 121 16 Prague, Czech Republic*

The possibility of using a quantum computer D-Wave 2X with more than 1000 qubits to determine the global minimum of the energy landscape of trained restricted Boltzmann machines is investigated. In order to overcome the problem of limited interconnectivity in the D-Wave architecture, the proposed RBM embedding combines multiple qubits to represent a particular RBM unit. The results for the lowest-energy (the ground state) and some of the higher-energy states found by the D-Wave 2X were compared with those of the classical simulated annealing (SA) algorithm. In many cases, the D-Wave machine successfully found the same RBM lowest-energy state as that found by SA. In some examples, the D-Wave machine returned a state corresponding to one of the higher-energy local minima found by SA. The inherently nonperfect embedding of the RBM into the Chimera lattice explored in this work (i.e., multiple qubits combined into a single RBM unit were found not to be guaranteed to be all aligned) and the existence of small, persistent biases in the D-Wave hardware may cause a discrepancy between the D-Wave and the SA results. In some of the investigated cases, introduction of a small bias field into the

**energy function or optimization of the chain-strength parameter in the D-Wave embedding successfully addressed difficulties of the particular RBM embedding. With further development of the D-Wave hardware, the approach will be suitable for much larger numbers of RBM units.**

## 1 Introduction

---

There is a strong expectation that adiabatic quantum computers (AQC) will benefit a wide range of applications, including machine learning and neural networks in particular (Farhi, Goldstone, Gutmann, & Sipser, 2000). The first commercial machine that is a quantum annealer (QA), the D-Wave (D-Wave Systems), was developed to determine the ground state (i.e., the global minimum) of systems with complex energy landscapes (complex Hamiltonians of a many-qubit system). (See Table 1 for a list of acronyms and select notation used in this letter.) An efficiency and reliability of the ground-state determination by D-Wave machines was extensively investigated recently (Boixo et al., 2014; Trummer & Koch, 2015, Novotny, Hobl, Hall, & Michielsen, 2016).

From an applications point of view, a QA is expected to speed up the training energy-based models such as restricted Boltzmann machines (RBM) (Smolensky, 1986; Bian, Chudak, Macready, & Rose, 2010; Fischer & Igel, 2014) and deep Boltzmann machines (Salakhutdinov & Hinton, 2009, 2012). For example, one of the most computationally demanding parts of the RBM training algorithm is sampling from the model distribution when computing the model-dependent term in the gradient of the log likelihood (Fischer & Igel, 2014). Ideally, it would require summing over all of the possible values of the RBM visible and hidden units. For this reason, this task is accomplished in practice by applying a Markov chain Monte Carlo (MCMC) method, which is computationally intensive and provides only an approximation to the desired sum.

Replacing some of the MCMC steps with algorithms that include a QA offers a promise of faster and more efficient training. The first reason for this promise is that the low-energy states found by QA have been demonstrated to follow a distribution close to the Boltzmann distribution (Rose, 2014; Adachi & Henderson, 2015; Benedetti, Reaple-Gómez, Biswas, & Perdomo-Ortiz, 2015). The expectation is that those states may be used to generate a proper set of samples from the model distribution very efficiently compared to MCMC algorithms. The second reason is that there is an expectation of achieving a more representative set of samples, compared, for example, to the results of Gibbs sampling. Benedetti et al. (2015) demonstrated that for a small RBM having 16 visible and 16 hidden units, their QA-based algorithm outperformed the conventional contrastive divergence (CD-1) and is capable of at least matching CD-100. However, before algorithms for bigger RBMs are investigated, it is important to know exactly what information

Table 1: Acronyms and Select Notation.

BAS	Bars and stripes patterns
GS	Ground state (the lowest-energy state)
MCMC	Markov chain Monte Carlo
RBM	Restricted Boltzmann machine
QA	Quantum annealer
SA	Simulated annealing
$E(\mathbf{v}, \mathbf{h})$	Energy function for the visible and hidden variables.
$Z$	Partition function
$L(\theta \mathbf{v}_{tr})$	Likelihood given the data $\mathbf{v}_{tr}$ and depending on parameters $\theta$
$\mathbf{s}$	State of Ising spins
$J_{ij}$	Weights of the D-Wave couplings between the two qubits $i$ and $j$
$h_j$	Bias field for a qubit $j$
$\mathbf{v}$	Vector of the visible nodes
$\mathbf{h}$	Vector of the hidden nodes
$b_j$	Bias of the visible RBM node $j$
$c_j$	Bias of the hidden RBM node $j$
$\omega_{ij}$	Weight between nodes $i$ and $j$

is obtained when a QA solves a problem specifically corresponding to the RBM energy function and how reliable this solution is for a wide range of possible RBM problems. Specifically, does the QA always return the true ground state of the RBM energy function, and what is the nature of any other higher-energy solutions returned?

For example, Novotny et al. (2016) investigated spanning tree problems that are considered easy for classical computers but difficult for the D-Wave due to the particularities of the problem's energy landscape. Only 42% of the spanning trees of different sizes successfully found the ground state at least once out of 100,000 quantum anneals. It was reported that some spanning trees do not find the ground state even after a million anneals. Spanning tree problems is rather unique and very different from those involving RBMs. Nevertheless, it is yet to be established whether the kind of the energy functions characteristic for RBMs can be reliably handled by a QA, a motivation for this work.

Some theoretical work has been done on predicting the applicability of the QA hardware to RBM problems (Dumoulin, Goodfellow, Courville, & Bengio, 2014). Since there is a limit of six couplings for each qubit in the current D-Wave hardware, the effects of the limited connectivity were investigated in that work. It had been established that RBM implementations with a progressively smaller number of connections between the units make it increasingly more difficult to learn the probability distribution represented by the training set. Later, the potential of the actual quantum hardware for training deep neural networks was investigated (Rose, 2014; Adachi et al., 2015; Benedetti et al., 2015). A few specific implementations of the neural network embedding and sampling from a D-Wave QA have been reported.

In this work, we empirically investigated what information can be obtained by applying a QA to the relatively rugged energy function of a fully trained RBM. Specifically, the focus was on investigating a particular RBM embedding within the Chimera lattice used in the D-Wave architecture, aiming at achieving a significant connectivity enhancement over the native Chimera RBM. Such enhancement was enabled by combining multiple qubits to represent a single visible or a single hidden RBM unit. The global minima returned by a D-Wave call for this particular RBM implementation were compared to the lowest-energy states found by simulated annealing (SA) performed on a classic computer.

In addition, probabilistic models trained with energy-based machine learning algorithms are expected to provide an additional verification that QA is capable of correctly determining the true ground state of systems with very rugged energy landscapes. By virtue of energy-based learning, the RBM training algorithm strives to minimize the energies (i.e., maximize all or some of the joint probabilities) corresponding to the marginal probabilities for the training patterns while maximizing the energies associated with all of the other states of the system.

Thus, while traditional classical computing algorithms such as SA cannot be relied on to verify the true global energy minimum reported by the D-Wave, the energy-based learning at least makes it increasingly probable that the GS value of the state vector should become one of the training patterns (see the discussion in section 3.1). Thereby, a combination of a comparison of D-Wave results with results of SA and the expectation that one of the training patterns is likely to emerge as the GS after the RBM training may offer additional support for the correctness of the D-Wave GS determination for complex energy landscapes.

Naturally, the described objective of this work does not adequately answer the main question about the possibility of efficient QA-based sampling from the model distribution of large-size RBMs. While approaches based on, for example, estimating the effective sampling temperature (Benedetti et al., 2015) or possibilities of simulated warming from the lowest-energy states are being investigated, our ability to trust the correctness of the QA ground state and the lowest excited states specifically for particular RBM architectures is an important prerequisite to developing those methods.

## 2 Adiabatic Quantum Computer

---

**2.1 Quantum Annealing.** One of the popular references for the adiabatic quantum computer (AQC) is Farhi et al. (2000), and a comprehensive overview of the AQC concept can be found in Santoro and Tosatti (2006). AQC employs quantum annealing to find a ground state (i.e., a global energy minimum) of a spin glass. It serves as an alternative to the classical SA method, which is based on thermal effects. In a QA, the search for the global minimum of the energy functional of the problem uses not only thermal

fluctuations over barriers of the classical SA algorithm, but also quantum mechanical tunneling through the barriers surrounding local minima. Full details of the functioning of an ideal AQC or how well the QA D-Wave machine functions are beyond the scope of this letter. Instead, the availability of a QA hardware enables one for the first time to examine if a QA algorithm can handle a relatively large trained Boltzmann machine, which is the scope of this work.

**2.2 QA D-Wave.** The D-Wave hardware specifically implements an Ising spin-glass model (Binder & Young, 1986; Lucas, 2014) using qubits. The underlying graph for the qubit interconnections for this problem is the so-called Chimera graph represented by a set of  $N$  vertices, which correspond to the optimization variables  $s_j$ , and couplings between some of the qubits at the vertices. The weights of the couplings  $J_{ij}$  and the bias fields  $h_j$  are defined to represent a particular spin glass problem for which the ground state is to be solved. The ground state of the Ising spin glass is the state  $\mathbf{s}$  of Ising spins that minimizes the energy function (Stein & Newman, 2013):

$$E(\mathbf{s}) = - \sum_{i=1}^{N-1} \sum_{j=1+1}^N J_{ij} s_i s_j - \sum_{j=1}^N h_j s_j \quad (2.1)$$

### 3 Restricted Boltzmann Machines

---

The architecture for the general Boltzmann machine (BM) is a particular case of the general Ising spin glass model. Similar to D-Wave's qubits, the RBM units accept binary values ( $\{0, 1\}$  or  $\{-1, 1\}$  depending on the implementation). In the following, we describe the Ising representation of the RBM ( $-1$  and  $1$  spin-down and spin-up values instead of the traditional in machine learning  $0$  and  $1$  values).

The weights of the connections  $\omega_{ij}$  between RBM units are similar to D-Wave's couplings  $J_{ij}$ , and the biases  $b_i$  and  $c_j$  for each RBM unit play roles similar to the local fields  $h_i$  for the D-Wave qubits. An important difference for machine learning applications is the existence of restrictions on the values of  $J_{ij}$  and  $h_i$  in the D-Wave hardware, which are limited to the range  $[-1, 1]$ . Most important, there is a difference between the D-Wave and a general RBM in the number of connections. For the D-Wave, the number of couplings for each qubit is limited to six (see Figure 1). Furthermore, some qubits have even smaller numbers of couplings, which is caused by either the particular physical coupling being inactive or an entire qubit missing in the specific implementation of the particular hardware chip.

The RBM units correspond to a complete bipartite graph and are divided into  $m$  visible units  $\mathbf{v} = \{v_1, v_2, \dots, v_m\}$ , which are used to represent the observations, and  $n$  hidden units  $\mathbf{h} = \{h_1, h_2, \dots, h_n\}$ , serving as the latent

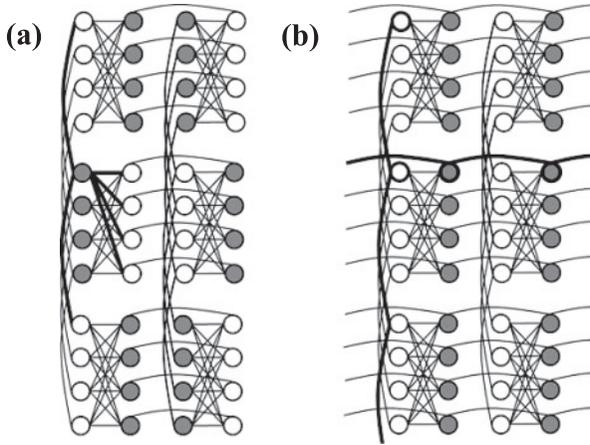


Figure 1: (a) A straightforward way of RBM embedding into the Chimera lattice, with each visible unit connected only to hidden units and vice versa. The open and filled circles are qubits corresponding to visible and hidden units, respectively. At most, six connections are available for each RBM unit. (b) One possible way of embedding with the objective of increasing the number of connections for each unit in a Chimera lattice. The bold lines highlight “ferromagnetic” couplings between multiple qubits, which combine them into a single logical RBM unit (the vertical bold lines combine visible and the bold horizontal lines combine hidden units).

variables. The visible and hidden units have biases  $b_j$  and  $c_i$ , respectively. Each  $v_i$  is connected to all  $h_j$  units with weight  $\omega_{ij}$  and vice versa. As a result, the full RBM implementation has the number of connections for each unit equal to the number of units from the other category. While usually the term *RBM* refers to the BM implementation having all possible connections between the visible and hidden units, in this letter, we use *RBM*, with proper additional notations, to also mean implementations where every visible unit has connections only to a subset of hidden units. Specifically, we refer to them as “Chimera-RBM” and “the second RBM embedding” (as introduced in sections 4.1 and 4.2).

The D-Wave qubits have a limited number of connections. For the native Chimera RBM implementation, for which each qubit represents a particular RBM unit (see section 4.1), there are at most six connections between a given visible unit and the hidden units. We found this to be insufficient for training RBMs of a reasonable size. This was evidenced by an inability to achieve classification errors exceeding 60%, which is also supported by the earlier published results on the effect of the limited connectivity (Dumoulin et al., 2014).

**3.1 RBM Energy Function.** The RBM's energy function  $E(\mathbf{v}, \mathbf{h})$ , is a particular case of the energy function of a general Ising spin glass model:

$$E(\mathbf{v}, \mathbf{h}) = - \sum_{i=1}^n \sum_{j=1}^m \omega_{ij} h_i v_j - \sum_{j=1}^m b_j v_j - \sum_{i=1}^n c_i h_i. \quad (3.1)$$

Consequently, RBM embedding into the Chimera lattice would allow using QA to find the low-energy states of  $E(\mathbf{v}, \mathbf{h})$ . The joint probability distribution of the model is the Gibbs distribution,

$$p(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} e^{-E(\mathbf{v}, \mathbf{h})/T}, \quad (3.2)$$

where the partition function is

$$Z(T) = \sum_{\mathbf{v}, \mathbf{h}} e^{-E(\mathbf{v}, \mathbf{h})/T},$$

and  $T$  is a temperature-like parameter. In most RBM machine learning applications, the training and inference are performed at  $T = 1$ . This is different from the D-Wave, which operates at a very low physical temperature. Specifically, while the physical temperature in the current hardware is about 15 mK, the effective temperature relevant for modeling the distribution of samples from QA with a Boltzmann distribution may be different and must be determined for the particular situation (Benedetti et al., 2015).

**3.2 RBM Training by Log-Likelihood Maximization.** Training of any energy-based model is based on maximizing the probability for the particular states corresponding to the training patterns in the training set. This can be accomplished by performing minimization of a loss function quantifying a difference between the model and the training probability distributions. In turn, this is accomplished by maximizing the marginal probability distribution  $p(\mathbf{v})$  for the desirable configurations  $\mathbf{v}_{\text{tr}}$ , while keeping it low for all the other possible configurations. It should be noted that while the maximum  $p(\mathbf{v})$  corresponding to  $\mathbf{v} = \mathbf{v}_{\text{tr}}$  could in general be consistent with the joint probability  $p(\mathbf{v}, \mathbf{h})$  having its maximum at a state with  $\mathbf{v} \neq \mathbf{v}_{\text{tr}}$ , such probability distribution would go against the trend of the entropy maximization (Jaynes, 1957) during the log-likelihood maximization. Therefore, the chance of achieving the ground state and the local minima represented by some of the training patterns must increase during the training, even though the maximum entropy that is consistent with the training objectives may not be large enough to guarantee such an outcome, especially for more restrictive training objectives. For RBM training, maximization of the

marginal probability distribution is accomplished by gradient ascent optimization of the model parameters  $\theta$  to maximize the log likelihood given the training data  $\ln L(\theta | \mathbf{v}_{tr}) = \ln p(\mathbf{v}_{tr} | \theta)$ .

For an RBM, the parameters  $\theta$  are  $\omega_{ij}$ ,  $b_j$ , and  $c_i$ . The corresponding expressions for the gradient of the log likelihood for a single training pattern  $\mathbf{v}_{tr}$  are

$$\begin{aligned} \frac{\partial \ln L(\theta | \mathbf{v}_{tr})}{\partial \omega_{ij}} &= [2p(H_i = 1 | \mathbf{v}_{tr}) - 1] v_j \\ &\quad - \sum_{\mathbf{v}} p(\mathbf{v}) [2p(H_i = 1 | \mathbf{v}) - 1] v_j, \end{aligned} \quad (3.3)$$

$$\frac{\partial \ln L(\theta | \mathbf{v})}{\partial b_j} = v_j - \sum_{\mathbf{v}} p(\mathbf{v}) v_j, \quad (3.4)$$

$$\frac{\partial \ln L(\theta | \mathbf{v})}{\partial c_i} = [2p(H_i = 1 | \mathbf{v}_{tr}) - 1] - \sum_{\mathbf{v}} p(\mathbf{v}) [2p(H_i = 1 | \mathbf{v}) - 1]. \quad (3.5)$$

**3.3 MCMC and Simulated Annealing.** The most computationally intensive parts of equations 3.3 to 3.5 are the second terms, which generally require summation over all possible values of the vector of the hidden units. While QA-based algorithms offer a promise of an efficient computation for this part of the RBM learning algorithm, it is beyond the scope of this work. Instead, the D-Wave machine was applied to perform QA on the Ising spin glass problem of an already trained RBM (see section 4.1). For RBM training in this work, a standard approximation technique, Gibbs sampling, was used, which is one of the implementations of the Markov chain Monte Carlo (MCMC) algorithm. In this algorithm, a training pattern is used as the initial value for the visible units in the Gibbs chain. The following expressions give the conditional probability of a single variable being equal to one. Formulas 3.6 to 3.8 are used to sequentially sample at  $T = 1$  the next value of the vector of the visible or hidden units in the MCMC until either the stationary distribution or a specified limit of steps was reached:

$$p(H_i = 1 | \mathbf{v}) = \sigma \left( 2 \left( \sum_{j=1}^m \omega_{ij} v_j + c_i \right) \right), \quad (3.6)$$

$$p(V_j = 1 | \mathbf{h}) = \sigma \left( 2 \left( \sum_{i=1}^m \omega_{ij} h_i + b_j \right) \right), \quad (3.7)$$

$$\sigma(x) = 1/(1 + e^{-x/T}). \quad (3.8)$$

The same formulas can be used when SA is applied in order to sample the states corresponding to the global and local minima.

When maximizing the log-likelihood  $\ln L(\theta | \mathbf{v}_{tr}) = \ln p(\mathbf{v}_{tr} | \theta)$  given the training data, the gradient ascent algorithm aims at maximizing the sum of joint probabilities corresponding to the training patterns, which for a single training vector  $\mathbf{v}_{tr}$  has the following form:

$$p(\mathbf{v}_{tr} | \theta) = \sum_{\mathbf{h}} p(\mathbf{v}_{tr}, \mathbf{h} | \theta) = \frac{1}{Z(\theta)} \sum_{\mathbf{h}} e^{-E(\mathbf{v}_{tr}, \mathbf{h} | \theta)}. \quad (3.9)$$

As a result, such training favors directions that may maximize individual joint probabilities  $p(\mathbf{v}_{tr}, \mathbf{h} | \theta)$  in the sum (i.e., minimize  $-E(p(\mathbf{v}_{tr}, \mathbf{h} | \theta))$ ) at the expense of other joint probabilities corresponding to states  $(\mathbf{v}, \mathbf{h})$  other than the training patterns. While it is far from guaranteed, a significant degree of  $L(\theta | \mathbf{v}_{tr})$  maximization makes it more likely that the global minimum of the energy function will turn out to be a state  $(\mathbf{v}, \mathbf{h})$  corresponding to one of the training patterns  $\mathbf{v}_{tr}$ , at least for the cases of less restrictive sets of constraints, the training objectives (see our discussion in section 3.2).

## 4 Methods

---

**4.1 Chimera-RBM.** The Chimera topology of the D-Wave is an  $L \times L$  lattice with each 8-qubit unit cell representing a complete bipartite graph  $K_{4,4}$ . The 1000 qubit D-Wave 2X has  $L$  equal to 12, which results in 1152 total qubits possible, of which 55 are inactive in the chip we used.

A section of a Chimera lattice, with a few unit cells having the  $4 \times 2$  arrangement, is shown in Figure 1. Each unit cell has eight qubits. The Chimera lattice is a bipartite graph. Hence, the native implementation of an RBM on the Chimera graph could use the bipartite nature to assign one Boltzmann machine unit to each qubit. This results in an RBM that is further restricted to the Chimera graph and in which visible units are connected to at most six hidden units, and vice versa.

A straightforward way of achieving this kind of embedding of the RBM units into the Chimera lattice is shown in Figure 1a. The open and filled circles represent qubits corresponding to visible and hidden units, respectively. Every visible unit is connected only to hidden units and vice versa. In this way, at most six connections are available for each RBM unit. The values of D-Wave's couplings  $J_{ij}$  are set to the corresponding values of connections  $\omega_{ij}$  between RBM units, and the values of D-Wave's local fields  $h_i$  are set to the corresponding values of the RBM biases  $b_i$  and  $c_j$ . The missing connections (all but the six available) are accounted for by setting corresponding  $J_{ij}$  and  $\omega_{ij}$  to zero in equations 2.1 and 3.1, respectively. Henceforth, this embedding will be referred to as Chimera-RBM.

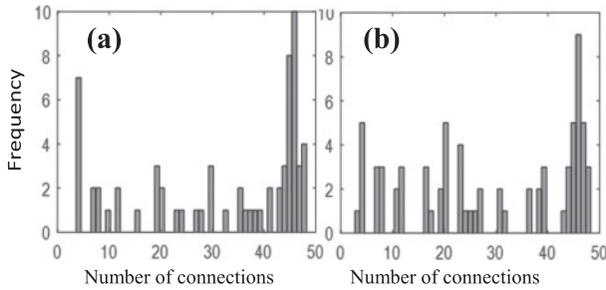


Figure 2: (a, b) Histograms of the available connections for each visible and hidden unit, respectively, for the D-Wave chip used in this study based on the embedding of Figure 1b.

**4.2 Connectivity Enhancement by Combining Qubits.** Figure 1b illustrates the second RBM embedding that was used in this work to increase the connectivity for RBM units beyond the limit of six connections available to individual qubits. The four qubits in the first column from the left of each  $4 \times 2$  cell are selected to serve as the visible units, and the four qubits in the second column serve as hidden units for our RBM embedding. By making the bonds corresponding to the bold horizontal lines in Figure 1b ferromagnetic (i.e., setting them to the maximum allowed value of one), multiple qubits from different cells are combined to try to align themselves in the same direction, thereby ensuring that multiple qubits represent a single hidden RBM unit. In a similar way, couplings corresponding to bold vertical lines in Figure 1b are used to combine multiple qubits to represent a single visible RBM unit.

If  $N$  qubits are combined to represent a particular RBM unit, this unit will have up to  $4N$  connections. The actual number may be smaller due to the missing couplings and qubits in the hardware. Applying this procedure for an ideal 1152-qubit Chimera lattice without any missing qubits or couplings would enable, for example, embedding a standard RBM with 48 visible and 48 hidden units, with the complete set of 48 connections for each visible or hidden unit.

The straightforward approach to achieve this embedding with the real D-Wave lattice having 55 missing qubits is to stop adding qubits to the particular RBM unit when a missing connection is encountered. This approach resulted in 66 visible and 73 hidden units, with fewer than 48 connections for each. The energy function is the same as in equation 3.1, but the missing connections are accounted for by setting corresponding  $\omega_{ij}$  to zero. The number of connections is different for different units, ranging from only 3 (for one of the hidden RBM units) to 48 (the histograms in Figures 2a and 2b). This embedding was achieved manually by setting the ferromagnetic bonds for the combined qubits to one and dividing the biases  $b_i$  or  $c_j$

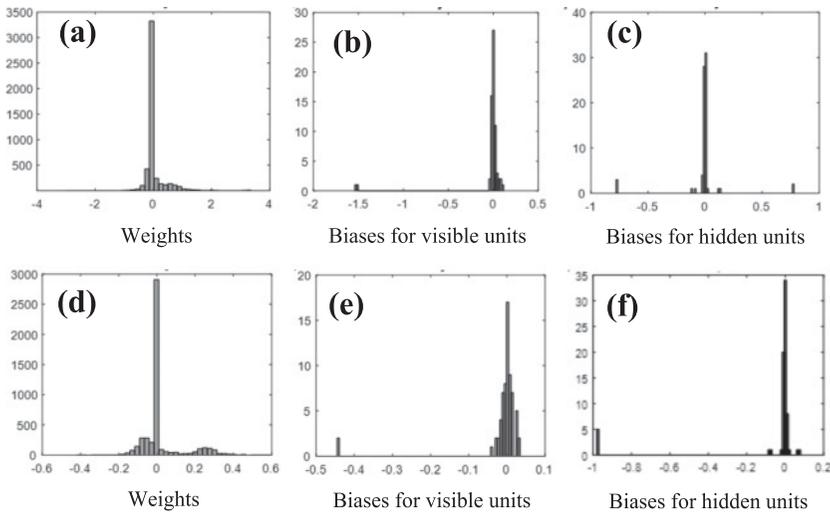


Figure 3: Histograms of (a) weights  $\omega_{ij}$ , (b) biases for the visible units  $b_i$ , and (c) biases for the hidden units  $c_j$  after 10,000 iterations of training RBM with 508 BAS training patterns without utilizing weight/bias decay. (c, d, e) The same with employing decay coefficients in order to keep the weights and biases within the range of  $[-1, 1]$ .

equally between the corresponding combined qubits. This approach is an approximation to a precise embedding using a constraint that the combined qubits representing a particular RBM unit all must have the same value (i.e., connected with ferromagnetic bonds). According to equation 3.1, when the ferromagnetic bond constraint is satisfied, the energy function of the embedding will have the same value as that of the original RBM if the biases of the combined qubits add up to the value of the bias of the corresponding RBM unit. Modest violations of the ferromagnetic bond constraint by the hardware make our approach an approximation to a precise embedding. In addition, a D-Wave’s subroutine `sapiEmbedProblem` (the references are not publicly available) was used for comparison, which, in addition, allowed varying the values of the so-called chain-strength parameter in search of better embedding conditions.

The second difficulty was the D-Wave’s requirement of keeping the values of the weights and biases in the  $[-1, 1]$  range. This requirement was satisfied by using weight-decay and bias-decay parameters (Fischer & Igel, 2014) in the training algorithm (see Figure 3). Finally, the Ising representation of the RBM ( $-1$  and  $1$  spin-down and spin-up values instead of the traditional in machine learning  $0$  and  $1$  values) was used for the ease of the data exchange with the D-Wave.

**4.3 RBM Training and Training Samples.** Investigation of the training of the described RBM model to achieve sufficient generalization for the pattern recognition and classification of the patterns previously unseen by RBM, for example, is beyond the scope of this work and will be reported elsewhere. The RBM training in this work was narrowly serving the goal of achieving a complex energy function profile with all the lowest-energy local minima related to the training patterns. Since the goal of this work was to investigate the D-Wave's capability to sample from the model probability distributions representative of typical RBM tasks, no precautions have been taken to prevent overtraining while producing a rough energy landscape for D-Wave experiments. The main outcomes of this work (a match between SA and QA results) were observed after partial training, as well as after training that results in the best classification error when using the classical computer.

For the given small RBM with connectivity increased by combining qubits, a toy problem of  $8 \times 8$  bars-and-stripes (BAS) was selected (MacKay, 2002; Hinton & Sejnowski, 1986). The training patterns are shown in the top rows of Figures 4 and 6 and the number of training patterns used in the particular experiment is indicated in the corresponding figure captions. The BAS problem was found to be attractive due to the ease of generating very diverse sets of training patterns, with a wide range of variation of the ratio of spin-up to spin-down unit values. This was also useful for verifying the possible existence of a slight unintentional bias field in the D-Wave. The BAS model also has a convenience that the training configurations are easily visualized. Two additional visible units were used as placeholders for the labels for the classification tasks to be pursued in the future work. However, the label values were not differentiated in this work (both labels were set to  $-1$  in all the training patterns used).

In addition to BAS, a set of 10 digits was generated; the results are shown in Figure 8. Another set of 10 smaller digits was generated. By shifting those digits in a different direction, 240 distinct patterns were achieved. Shifting and also inverting each of the 10 digits produced 480 distinct patterns for training.

**4.4 Determination of the Minima of the RBM Energy Function.** In this work, the D-Wave was run to generate the lowest-energy solutions for the embedding (described in section 4.2) of the fully trained RBM 1000 times in each D-Wave call, and the lowest-energy solutions were selected as the GS and the lowest excited states. SA was conducted on the actual RBM graph (i.e., the bipartite graph of 66 visible and 73 hidden units having the limited connectivity of fewer than 48 connections for each as described in section 4.2, but free from the unit's representation by combining qubits) on a classical computer to try to find independently local, and perhaps global, minima of the actual model distribution. The number of D-Wave anneals was 10,000. During SA, an initial state (i.e., the vector of the visible units) was

generated randomly, and a specified number of Gibbs sampling steps were conducted at an initial elevated temperature of  $T = 10$ . Following that, the value of  $T$  was reduced linearly in a specified number of temperature steps (typically, 1000 steps). In each temperature step, up to 1000 Monte Carlo sweeps were performed before moving to the next lower temperature. The algorithm was terminated when the value of  $T$  approached zero, making the temperature fluctuations around the minimum insignificant. The procedure was repeated using different random initial states, until no new low-energy states were generated at the end of SA. The solution with the lowest energy was chosen as the candidate for the ground state.

## 5 Results and Discussion

---

**5.1 Chimera-RBM.** A few different sizes of Chimera-RBM having different numbers of visible and hidden nodes were used for training with  $4 \times 4$ ,  $5 \times 5$ ,  $6 \times 6$ , and  $8 \times 8$  BAS training sets. A subset of the D-Wave qubits was selected, with each representing a particular RBM unit. Other qubits were excluded by setting their weights to zero. As discussed above, at most six connections are available to each unit.

Our observations qualitatively matched the results of Dumoulin et al. (2014). For the 16 visible units of the smallest RBM we attempted, the six available connections constituted more than a third of the possible connections of the full RBM. But even for this very small RBM, training was difficult and resulted in a very low percentage of correct recognitions after training, even when the set of the test patterns included only the same patterns that were used for training.

For RBMs with a larger number of visible units, it was found impossible to learn the probability distribution represented by the training set, which was caused by the fact that the number of connections for each unit of a bigger RBM represented a small fraction of the number of connections used in a fully connected RBM. Since much larger RBM sizes are required for practical applications, no training optimization for the  $4 \times 4$  BAS case was attempted, and the Chimera-RBM was not investigated any further.

**5.2 Small Number of Training Patterns.** The remainder of this letter describes the results for the second RBM embedding (see Figure 1b), in which the number of available connections was increased beyond the limit of six by combining multiple qubits to represent each RBM unit.

At first, the RBM having 66 visible and 73 hidden units was trained with a small number of patterns in an attempt to generate the energy function having a small predictable set of the local minima. Figure 4 shows the D-Wave GSs (which is the state corresponding to  $\arg \max_{\mathbf{v}, \mathbf{h}} p(\mathbf{v}, \mathbf{h})$ ) determined for RBMs trained with (a) one pattern, (b) two patterns, and (c) four patterns, all randomly selected. In this and other figures, each grid line corresponds to a pixel row or column. An individual separate pixel would show up as a

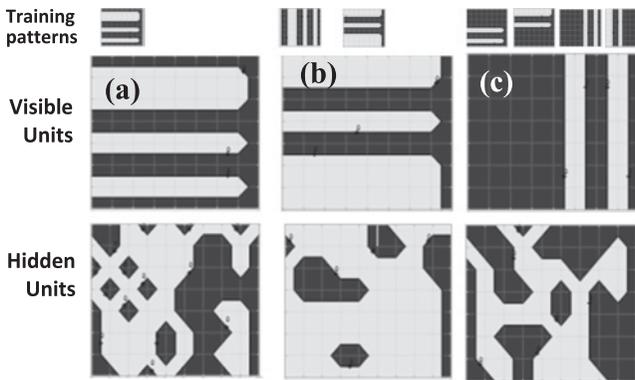


Figure 4: The lowest-energy states for the visible (top) and the hidden (bottom) units for RBMs trained with (a) one pattern, (b) two patterns, and (c) four patterns, all randomly selected. The dark and the light colors of the map respectively represent  $-1$  and  $1$  values of the corresponding RBM unit. The corresponding training patterns for each case are shown above each of the states. The same lowest-energy states were found by both the simulated annealing and the D-Wave.

rhomb. SA was performed on the trained RBM to find all, or at least some, of the low-energy local minima for the corresponding energy function. All of the lower-energy states found by SA in each of the three cases corresponded to one, two, and four training patterns, respectively. The visible and hidden units for the D-Wave GSs shown in Figure 4 exactly matched those for the lowest-energy state found by SA in each case.

We also examined the nature of other (higher-energy) solutions returned by the D-Wave machine. The second- and third-lowest-energy D-Wave states are shown in Figure 5 for the RBM trained with four patterns. The second D-Wave state (see Figure 5a) was used as the initial state for the energy gradient descent calculations (i.e., an MCMC conducted at zero temperature), and it quickly relaxed to the lowest-energy state. This behavior allowed us to establish that it was not a local minimum. The third D-Wave state happened to be one of the three remaining training patterns (see Figure 5b). This third D-Wave state was a local minimum of the RBM energy function, and it corresponded to the second-lowest-energy local minimum found by SA. Most of the other higher-energy states returned by the D-Wave for the examples of Figure 4, as well as for other RBM examples investigated in this work, were similar: states having increasingly higher Hamming distance from the GS. However, occasionally, some of the states corresponded to a local minimum (often a local minimum with visible units representing a training pattern other than that of the GS; see Figure 5b).

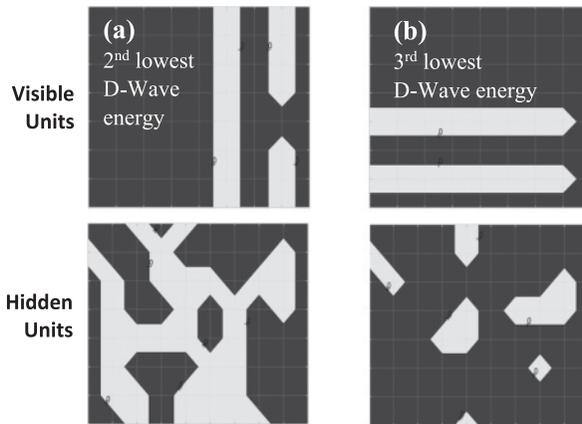


Figure 5: The second- and the third-lowest-energy states found by D-Wave for RBM trained with four patterns (see Figure 4c). The third D-Wave state happened to be a local minimum of the RBM energy function corresponding to the second-lowest-energy found by SA. The second D-Wave state is not a local minimum. During gradient descent in energy, it quickly relaxes to the lowest-energy state.

**5.3 Discrepancies between D-Wave and SA.** In some cases, the D-Wave had difficulty reliably determining the same lowest-energy state that was found by SA. For example, it happened when the D-Wave was applied to an RBM trained with 10 patterns (see Figure 6). The lowest-energy state returned by the D-Wave corresponded not to the lowest but to the second-lowest-energy state determined by SA. This problem was also occasionally encountered in other experiments of this work. In general, it could be explained by nonperfect correspondence between the RBM and its Chimera embedding. Specifically, setting the weights for selected qubits to one does not guarantee that all those qubits will end up aligned in the same direction (it is impossible to guarantee that those bonds are perfectly ferromagnetic). In other words, the energy function of the Chimera embedding is an approximation of the corresponding RBM energy function. In many cases, the “correct” lowest-energy state could be found by experimenting with the chain-strength parameter used with the embedding subroutine provided by the D-Wave. It is important to note, however, that the problem seemed to be observed only when the energy gap between the ground state and the first excited state is very small (which, however, is a typical result for an RBM trained with a large number of the training patterns).

**5.4 Role of a Small Bias Field in the D-Wave Hardware.** The occurrences when the D-Wave failed to return the same lowest-energy state as

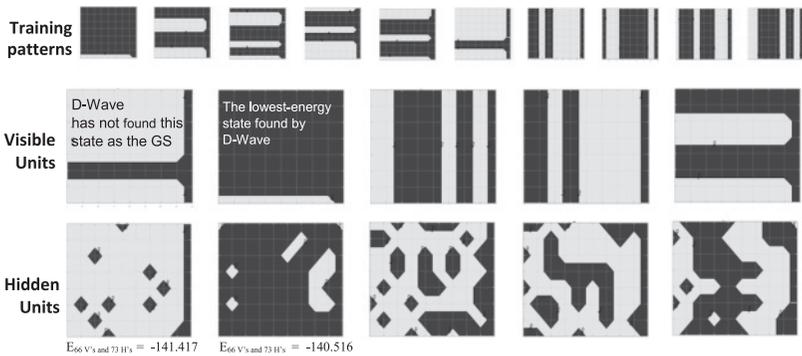


Figure 6: Five-lowest energy states found by SA for RBMs trained with 10 patterns (left to right from the lowest to higher energy). The 10 training patterns are shown in the top row. It was observed that a few lowest-energy local minima of the energy function after training corresponded to one of the training patterns. The lowest-energy state found by the D-Wave is not the first but the second state from the left, which is the second-lowest-energy state determined by SA. All the other D-Wave solutions represented states having increasingly higher Hamming distance from the GS.

the one found by SA often happened when the state missed by the D-Wave had spin-up qubits dominating the distribution. Preliminary evidence of the existence of a slight bias field favoring spin-down states has been reported (Novotny et al., 2016).

To verify if this mechanism could be responsible for the D-Wave missing the lowest-energy SA state as seen in Figure 6a, the RBM was trained with only two patterns, one of which was an inverted image (including the label units) of the other (see Figure 7). In this way, spin-up states dominated the pattern of Figure 7a, while the opposite was true for the inverted pattern of Figure 7b. After the RBM training, the two energy minima were found to correspond to two training patterns, with rather close values of the energy. The number of training iterations was varied to achieve the desired outcome—a state from the trained model distribution with visible units corresponding to pattern (a) had slightly lower energy than pattern (b), and it turned out to be the lowest-energy state determined by SA.

Similar to the case of Figure 6, the D-Wave missed the state of Figure 7a and returned as the GS a state having visible units of pattern b. In order to verify that a slight bias field in the hardware could be responsible for this GS identification failure, a small negative bias of  $-0.001$  was added to every qubit before the problem was sent to the D-Wave. Introduction of this artificial bias resulted in the D-Wave returning the same lowest-energy solution as SA (the state of Figure 7a). The second-lowest-energy state found by the D-Wave was pattern b. The other higher-energy states were states

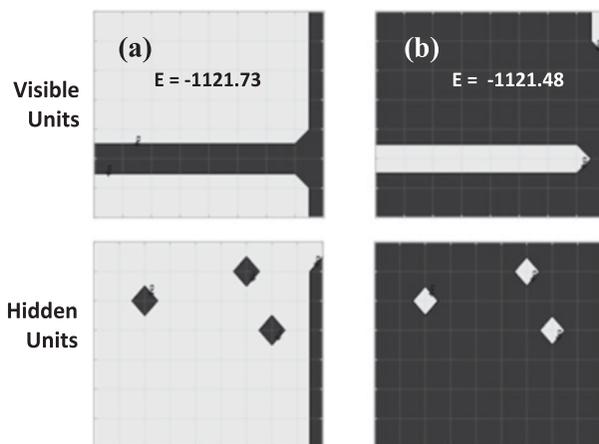


Figure 7: SA results for an RBM trained with only two patterns, the pattern in panel b being an inverted image of panel a. The RBM training was continued until the two minima had rather close values of energies. The lowest-energy state found by D-Wave was panel b. The RBM’s lowest-energy state found by SA annealing was panel a. Introduction of an additional artificial negative bias of  $-0.001$  resulted in D-Wave returning, in panel a, as the lowest-energy state.

with increasingly higher Hamming distance from that of Figure 7b, with progressively decreasing numbers of occurrences.

The existence of persistent, systematic biases in all of the available programmable parameters  $\{h_i, J_{ij}\}$  of the QA, which may be different for different qubits, was indeed recently demonstrated, and a procedure to quantify those biases was proposed (Perdomo-Ortiz, O’Gorman, Fluegemann, Biswas, & Smelyanskiy, 2015).

**5.5 The Full Set of the Training BAS and Digits Patterns.** Before investigating large training sets, a different kind of training set with 10 training patterns was implemented. Figure 8 reports an experiment similar to Figure 6 but with the 10 BAS patterns replaced with 10 digits. No problems were encountered. The GS found by the D-Wave was the same state as the first-lowest-energy state determined by SA.

Observations for higher-energy solutions returned by the D-Wave were qualitatively similar to the BAS examples described above. The second-lowest-energy state found by the D-Wave happened to be the state of the digit 4 (see Figure 8c), which is the third-lowest-energy minimum found by SA. Other D-Wave solutions represented nonlocal minima states having increasingly higher Hamming distance from the GS (see Figure 8a) or local minima corresponding to some other training patterns.

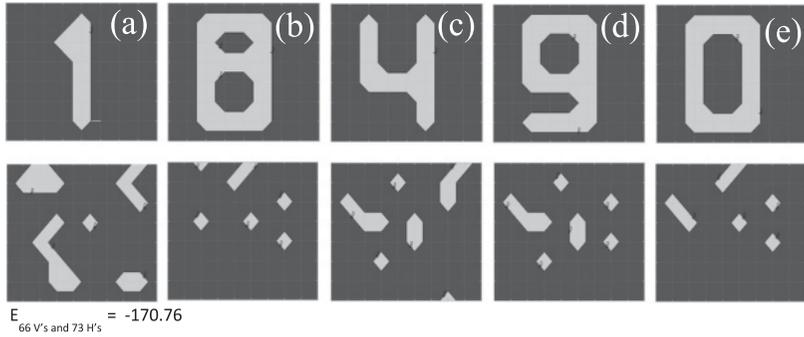


Figure 8: Five lowest-energy states found by SA for an RBM trained with 10 patterns representing digits from 0 to 9 (left to right from the lowest to higher energy). The GS found by the D-Wave was the state, panel a, which is the lowest energy state determined by SA. The second-lowest-energy D-Wave's state happened to be the state of the digit 4, panel c. Other D-Wave solutions represented nonlocal minima states having increasingly larger Hamming distance from the GS (a) or local minima corresponding to some other training patterns.

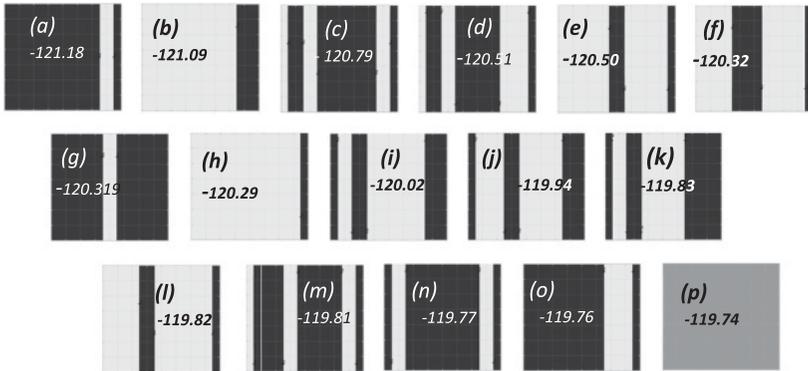


Figure 9: SA results for an RBM trained with all 508 BAS combinations. Only the visible units for the 16 lowest-energy states are shown. These energies for each state were calculated for the RBM with 66 visible and 73 hidden units.

Next, an RBM was trained with all possible 508 training BAS patterns. Sixteen lowest-energy local minima determined by SA, with the corresponding values of energies, are shown in Figure 9. In this experiment, the D-Wave was successful in determining the same lowest-energy state of the visible units as the one determined by SA annealing (see Figure 9a). Of course, for the energy function with such a large number of local minima, there is no guarantee that SA itself determined the correct ground state. At

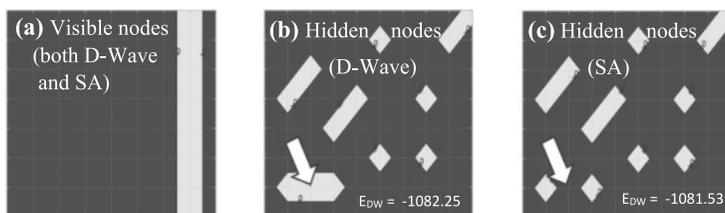


Figure 10: The lowest-energy state for RBMs trained with all 508 BAS combinations (pattern a in Figure 9): (a) the visible unit, which is the same for the D-Wave and for SA; (b, c) the corresponding hidden units found by D-Wave and SA, respectively. There is a Hamming distance of 1 between the vectors of hidden units found by the D-Wave versus SA. This is due to the violation of the ferromagnetic bond (aligned qubits) for some of the qubits corresponding to a particular hidden unit, which makes the D-Wave energy a bit lower than what is expected for the perfect ferromagnetic coupling of the combined qubits.

least, the fact that the visible units of the GS corresponded to one of the 508 training patterns and the fact that both the D-Wave and SA found the same lowest-energy state additionally increase the likelihood that the GS was found.

There was a small difference between the vectors of the hidden units corresponding to the lowest-energy solutions determined by the D-Wave and SA. It should be noted that in cases when not all the qubits combined to correspond to a particular unit had the same value, the value of the given unit was selected using the majority vote among the combined qubits. Figures 10b and 10c reveal the existence of a small Hamming distance of 1 between the hidden units found by the D-Wave versus SA, both corresponding to the identical vectors of the visible units. It was found to be again due to a violation of the ferromagnetic bond requirement for some of the qubits corresponding to a particular hidden unit (i.e., not all qubits that were intended to represent a given RBM unit were aligned by the QA in the same direction). This made the energy of this particular D-Wave state in Figure 10b a bit lower (i.e., more favorable) than the energy for the state having perfect ferromagnetic coupling of the combined qubits (which is the D-Wave state that would exactly correspond to the GS of the trained RBM).

Finally, instead of using BASs as the training sets, 240 as well as 480 digits were used. The 240 patterns were generated by shifting basic patterns that were similar to but somewhat smaller than digits shown in Figure 8. The set of 480 training patterns was generated by additionally inverting each of the 240 patterns. The D-Wave performance in those experiments was more problematic. A violation of the ferromagnetic bond requirement was found often and happened for many qubits corresponding to hidden units. As a result, the GS returned by the D-Wave corresponded to a state different from

the lowest-energy state found by SA. The visible units of the D-Wave GS matched one of the local minima found by SA. However, the visible units corresponded not to the lowest but to one of the higher-energy states found by SA.

Adjustments of the chain-strength parameter employed in the D-Wave's `sapiEmbedProblem` subroutine in a wide range of values did not cause desirable improvements. Alternatively, larger values of the weight decay parameters were used in an attempt to ensure that all of the RBM weights were significantly smaller in magnitude than the  $-1$  weights used for combining qubits into a single RBM unit. For the training cases of 240 and 480 digits, this approach was not successful. At a certain point, too much of a constraint on the magnitude of the weights makes it impossible to achieve sufficient RBM training.

## 6 Conclusion

---

The possibility of employing QAs to efficiently determine the GS or the lowest-energy excited states of a large RBM is potentially attractive for machine learning applications. Quickly determined lowest-energy states could be used, for example, to speed up generation of actual samples from the model distribution. Our ability to trust the correctness of the QA-generated solutions for particular RBM architectures is an important prerequisite to developing those methods.

We investigated the low-energy configurations returned by a D-Wave machine with more than 1000 qubits. The  $K_{4,4}$  Chimera lattice of the D-Wave has  $L \times L$  unit cells of 8 qubits each. A Chimera lattice without missing qubits would enable a complete RBM with  $4L$  visible and  $4L$  hidden units. Because some qubits are missing, we have mapped a further restricted Boltzmann machine onto the available D-Wave Chimera lattice chip. Further development of the QA hardware should make it possible to apply a similar algorithm to much larger numbers of RBM units.

We considered a variety of training cases, which provided an increasingly more complex energy landscape by increasing the number (and, later, the complexity) of the training patterns. In many of the investigated cases, the D-Wave found the same RBM lowest-energy state as that found by SA. While there was no guarantee that SA itself had not missed the true GS, in many experiments, this concern was alleviated by the property of the energy-based learning to give the lowest energies to the states having the values of the vector of all visible units corresponding to one of the training patterns. At least in the cases when SA found every single pattern that had been used for training, the match of the D-Wave and SA results is a rather reliable reassurance that the D-Wave found the true GS.

In some cases, some of the higher-energy solutions (i.e., the states with energy above the GS) found by the D-Wave corresponded to the second, third, or other local minima. It remains to be determined to what extent

such information about D-Wave's higher-energy states could be of potential interest for using QA to obtain multiple samples from the model, which has been explored, for example, in Rose (2014) and Adachi & Henderson (2015).

In a few experiments, the D-Wave returned as the GS, a state of visible units corresponding not to the lowest but the second- or sometimes a higher lowest-energy local minimum found by SA. This appears to be more likely to happen when the energy of different RBM local minima had very close values. In many cases, the problem can be eliminated by introducing a small bias field into the energy function. In other cases, the possibility of a non-perfect embedding of the RBM into the Chimera lattice (i.e., multiple qubits combined with close to but not always perfectly ferromagnetic bonds to enhance connectivity) is responsible for this effect and could be fixed in many (though not all) cases by experimenting with the chain-strength parameter used in the D-Wave embedding.

The BAS patterns investigated in this work do not necessarily result in the most rugged and challenging energy landscape. The D-Wave performance was further tested using more complicated training sets containing 10, 240, and 480 digits. While the 10-digits case did not cause any difficulties, the training cases with large numbers of digits consistently caused the noted problems of the ferromagnetic bond violation, which resulted in the D-Wave failing to find the RBM state corresponding to the global minimum found by SA. This difficulty is with the non-perfect embedding, in particular, the chaining together of physical qubits into a virtual RBM unit. Before the D-Wave becomes a useful tool for efficient sampling from the model distribution of large-size RBMs, possibilities of better embedding must be further investigated.

## Acknowledgments

---

Funding for this work was partially provided by the Pacific Northwest National Laboratory, under U.S. Department of Energy contract DE-AC05-76RL01830. In addition, part of Y.K.'s time was funded by the U.S. Army Research Office grant W911NF-14-1-0673 to support the U.S. Army RDECOM CERDEC NVESD. We acknowledge the support of the Universities Space Research Association Quantum AI Lab Research Opportunity Program.

## References

---

- Adachi, S. H., & Henderson, M. P. (2015). *Application of quantum annealing to training of deep neural networks*. arXiv:1510.06356
- Benedetti, M., Reaple-Gómez, J., Biswas, R., & Perdomo-Ortiz, A. (2015). *Estimation of effective temperatures in a quantum annealer and its impact in sampling applications: A case study towards deep learning applications*. arXiv:1510.07611v2

- Bian, Z., Chudak, F., Macready, W. G., & Rose, G. (2010). *The Ising model: Teaching an old problem new tricks*. [http://www.dwavesys.com/sites/default/files/weightedmaxsat\\_v2.pdf](http://www.dwavesys.com/sites/default/files/weightedmaxsat_v2.pdf).
- Binder, K., & Young, A. P. (1986). Spin glasses: Experimental facts, theoretical concepts and open questions *Rev. Mod. Phys.* 58, 801.
- Boixo, S., T. F. Rønnow, T. F. Isakov, S. V., Wang, Z., Wecker, D., Lidar, D. A., . . . Troyer, M. (2014). Evidence for quantum annealing with more than one hundred qubits *Nature Physics*, 10, 218–224.
- Dumoulin, V., Goodfellow, I. J., Courville, A. C., & Bengio, Y. (2014). On the challenges of physical implementations of RBMs In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence* (pp. 1199–1205). Cambridge, MA: MIT Press.
- Farhi, E., Goldstone, J., Gutmann, S., Sipser, M. (2000). *Quantum computation by adiabatic evolution*. arXiv:quant-ph/0001106
- Fischer, A., & Igel, C. (2014). Training restricted Boltzmann machines: An introduction. *Pattern Recognition*, 47, (1), 25–39.
- Hinton, G. E., & Sejnowski, T. J. (1986). Learning and relearning in Boltzmann machines. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*, vol. 1, *Foundations* (pp. 282–317). Cambridge, MA: MIT Press.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106, 620–630.
- Lucas, S. A. (2014). Ising formulations of many NP problems. *Front. Physics*, 12, 5.
- MacKay, D. J. C. (2002). *Information theory, inference and learning algorithms*. Cambridge: Cambridge University Press.
- Novotny, M. A., Hobl, L., Hall, J. S., & Michielsen, J. S. (2016). Spanning tree calculations on d-Wave 2 machines. *Journal of Physics: Conference Series*, 681, 012005. <https://iopscience.iop.org/article/10.1088/1742-6596/681/1/012005/meta>.
- Perdomo-Ortiz, A., O’Gorman, B., Fluegemann, J., Biswas, R., & Smelyanskiy, V. N. (2015). *Determination and correction of persistent biases in quantum annealers*. arXiv:1503.05679v1
- Rose, G. (2014). *First ever DBM trained using a quantum computer*. <https://dwave.wordpress.com/2014/01/06/first-ever-dbm-trained-using-a-quantum-computer/>.
- Salakhutdinov, R. R., & Hinton, G. E. (2009). Deep Boltzmann machines. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics* (vol. 12). JMLR.
- Salakhutdinov R. R., & Hinton G. E. (2012). An efficient learning procedure for deep Boltzmann machines. *Neural Comput.*, 24(8), 1967–2006. doi:10.1162/NECO\_a\_00311
- Santoro G. E., & Tosatti, E. (2006). Topical review: Optimization using quantum mechanics: Quantum annealing through adiabatic evolution *J. Phys. A: Math. Gen.*, 39, R393–R431.
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*, vol. 1: *Foundations* (pp. 194–281). Cambridge, MA: MIT Press.

- Stein, D. L., & Newman, C. M. (2013). *Spin glasses and complexity*. Princeton, NJ: Princeton University Press.
- Trummer, I., & Koch C. (2015). *Multiple query optimization on the D-Wave 2X Adiabatic Quantum Computer*. arXiv:1510.06437

---

Received November 2, 2016; accepted February 16, 2017.