



Published in final edited form as:

npj Quantum Inf. 2018 ; 4: . doi:10.1038/s41534-018-0060-8.

Quantum annealing versus classical machine learning applied to a simplified computational biology problem

Richard Y. Li^{1,2,3}, Rosa Di Felice^{2,4,5}, Remo Rohs^{1,2,4,6}, and Daniel A. Lidar^{1,3,4,7}

¹Department of Chemistry, University of Southern California, Los Angeles, California 90089, USA

²Computational Biology and Bioinformatics Program, Department of Biological Sciences, University of Southern California, Los Angeles, CA 90089, USA

³Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California 90089, USA

⁴Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA

⁵Center S3, CNR Institute of Nanoscience, Via Campi 213/A, 41125 Modena, Italy

⁶Department of Computer Science, University of Southern California, Los Angeles, CA 90089, USA

⁷Department of Electrical Engineering, University of Southern California, Los Angeles, California 90089, USA

Abstract

Transcription factors regulate gene expression, but how these proteins recognize and specifically bind to their DNA targets is still debated. Machine learning models are effective means to reveal interaction mechanisms. Here we studied the ability of a quantum machine learning approach to predict binding specificity. Using simplified datasets of a small number of DNA sequences derived from actual binding affinity experiments, we trained a commercially available quantum annealer to classify and rank transcription factor binding. The results were compared to state-of-the-art classical approaches for the same simplified datasets, including simulated annealing, simulated quantum annealing, multiple linear regression, LASSO, and extreme gradient boosting. Despite technological limitations, we find a slight advantage in classification performance and nearly equal ranking performance using the quantum annealer for these fairly small training data sets. Thus, we propose that quantum annealing might be an effective method to implement machine learning for certain computational biology problems.

Correspondence and requests for materials should be addressed to R.R. (rohs@usc.edu) or D.A.L. (lidar@usc.edu).

COMPETING INTERESTS

The authors declare no competing financial interests.

AUTHOR CONTRIBUTIONS

R.L., R.D.F., R.R. and D.L. designed and conceived the study. R.L. implemented and executed DW, SA, SQA, MLR, Lasso and XGB, and analyzed the results. R.R. and D.L. directed the study. R.L., R.D.F., R.R., and D.L. wrote the manuscript.

INTRODUCTION

Quantum computation has been the subject of intense scientific scrutiny for its potential to solve certain fundamental problems, such as factoring of integers [1] or simulation of quantum systems [2, 3], more efficiently than classical algorithms, by using unique quantum phenomena including entanglement and tunneling. More recently, there has been much interest in the potential of quantum machine learning to outperform its classical counterparts [4–18]. Although different implementations and models of quantum computing are still in development, promising theoretical and experimental research indicates that quantum annealing (QA) [19], or adiabatic quantum optimization [20], may be capable of providing advantages in solving classically-hard problems that are of practical interest (for a review see Ref. [21]). QA is the only paradigm of quantum computation that currently offers physical implementations of a non-trivial size, namely the D-Wave processors [22–24].

The adiabatic theorem of quantum mechanics, which underlies QA, implies that a physical system will remain in the ground state if a given perturbation acts slowly enough and if there is a gap between the ground state and the rest of the system's energy spectrum [25] (Fig. 1a). To use the adiabatic theorem to solve optimization problems, we can specify an initial Hamiltonian, H_B , whose ground state is easy to find (typically a transverse field), and a problem Hamiltonian, H_P , that does not commute with H_B and whose ground state encodes the solution to the problem we are seeking to optimize [26]. We then interpolate from H_B to H_P by defining the combined Hamiltonian $H(s) = A(s)H_B + B(s)H_P$, with $0 \leq s \leq 1$, where $A(s)$ and $B(s)$ are, respectively, decreasing and increasing smoothly and monotonically, t is time, and t_f is the total evolution, or annealing time. The adiabatic theorem ensures that the ground state of the system at $t = t_f$ will give the desired solution to the problem, provided the interpolation is sufficiently slow, i.e., t_f is large compared to the timescale set by the inverse of the smallest ground state gap of $H(s)$ and by $dH(s)/ds$ [27] (Fig. 1a). When QA is implemented in a physical device, temperature and other noise effects play an important role; thermal excitation and relaxation cannot be neglected and affect performance [28–30].

Quantum annealing algorithms were implemented on the D-Wave Two X (DW2X) processor installed at the Information Sciences Institute of the University of Southern California. The problem Hamiltonians that are used for D-Wave (DW) can be described as Ising spin models with tunable parameters [31]. The Ising model assumes a graph $G = (V, E)$ composed of a set of vertices, V , and edges, E . Each of the N spins is a binary variable located at a unique vertex. For the DW2X, $N = 1098$, the spins are represented by superconducting flux qubits, and G is the so-called Chimera graph (see Supplementary Material, Fig. S1). The problem, or Ising, Hamiltonian for this system can be written as

$$H_P = \sum_{i \in V} h_i \sigma_i^z + \sum_{(i,j) \in E} J_{ij} \sigma_i^z \sigma_j^z, \quad (1)$$

where the local fields $\{h_i\}$ and couplings $\{J_{ij}\}$ define a problem instance, and are programmable on the DW2X to within a few percent Gaussian distributed error. The σ_i^z

represent both binary variables taking on values ± 1 , and the Pauli z -matrices. Given a spin configuration $\{\sigma_i^z\}$, H_P is the total energy of the system. The initial Hamiltonian is a transverse magnetic field: $H_B = \sum_i \sigma_i^x$, where σ_i^x is the Pauli x -matrix acting on qubit i . The Ising Hamiltonian can be easily transformed into a quadratic unconstrained binary optimization (QUBO) problem by applying the following transformation:

$$w_i = \frac{\sigma_i^z + 1}{2} \in \{0, 1\}.$$

Problems submitted to DW are automatically scaled so that all h_i and J_{ij} values lie between -1 and 1 , and DW returns a set of spin values $\{\sigma_i^z = \pm 1\}$ that attempts to minimize the energy given by Eq. (1) (a lower energy indicates better optimization). Much attention has been paid to whether the DW devices are capable of delivering quantum speedups [32–36]. Here we sidestep this question and instead use the DW2X as a physical device that implements quantum annealing for the purpose of solving a problem in machine learning, while focusing on performance measures other than speedup.

In order to probe the potential of a machine learning approach that is based on quantum annealing, we have used the DW2X processor to solve the simplified formulation of a biologically relevant problem: transcription factor (TF)-DNA binding (see Fig. 1b). TFs are a key component in the regulation of gene expression, yet the mechanisms by which they recognize their functional binding sites in a cell and thereby activate or repress transcription of target genes are incompletely understood. Nucleotide sequence, flexibility of both TFs and binding sites, the presence of cofactors, cooperativity, and chromatin accessibility are all hallmarks that affect the binding specificity of TFs in vivo [37, 38]. As a first step to gaining insight into TF binding, it is valuable to understand the intrinsic binding specificity of the TFs for DNA, which is optimally gained from in vitro data. Widely used methods to gain such an understanding and represent the DNA sequence preferences of TFs are based on position weight matrices (PWM) or PWM-like models [39]. In the simplest of these models, the binding preference of a TF for each of the four nucleotides of the DNA alphabet {A, C, G, T} of a sequence of length L is represented as a $4 \times L$ matrix. Such models implicitly treat each position in the DNA sequence as being independent, so that each element of the matrix can be thought of as the contribution of a nucleotide at the corresponding position to the overall binding affinity. Since the independence of the nucleotide positions is in many cases a valid approximation and also because of current restrictions on the size of the DW processors, in this work we have used a model consisting of single-nucleotide sequence features to show a proof of principle of the use of machine learning via quantum annealing in biology. Despite technological limitations of emerging quantum technology, we concurrently demonstrate cases in which this form of machine learning using quantum annealing outperforms classical machine learning when training with small datasets. This is among the very first successful applications of quantum hardware to a realistic, though simplified problem in computational biology.

RESULTS

Experimental datasets on TF-DNA binding for a specific TF consist of N sequences of fixed length L and N values that express a measure of the binding affinity of the chosen TF to each

sequence: $\{(\vec{x}_n, y_n)\}_{n=1}^N$. In other words, the n th sequence is represented by the vector $\vec{x}_n = (x_{n,1}, x_{n,2}, \dots, x_{n,L})$ with $x_{n,j} \in \{A, C, G, T\}$, for $j = 1, \dots, L$, and y_n is the corresponding measure of binding affinity. For instance, \vec{x}_n may be ACAACTAA, with $y_n = 4.95$. In this work we used binding from three genomic-context protein binding microarray (gcPBM) experiments, which use fluorescence intensity as a measure of binding affinity, [40] and two high-throughput systematic evolution of ligands by exponential enrichment (HT-SELEX) [41–43] experiments, which report relative binding affinity. After preprocessing, the three gcPBM datasets consisted of $N \approx 1600$ sequences of $L = 10$ base-pairs. The two HT-SELEX datasets consisted of $N \approx 3200$ and 1800 sequences of length $L = 12$ after preprocessing (see Methods for a brief description of the preprocessing procedure). We used the following one-hot encoding to represent the sequence as a vector of binary variables: A = 1000, C = 0100, G = 0010, T = 0001, and thus transformed \vec{x}_n into a feature vector $\vec{\phi}_n \equiv (\phi_{n,1}, \dots, \phi_{n,4L})^T$. This encoding scheme [44] was used so that all combinations of inclusion and exclusion of the four nucleotides may be identified. Similar to previous studies [44–47] the goal of the present work is to identify patterns within the data to qualitatively assess whether the strength of a TF binding to a particular unseen sequence is above a certain threshold (classification) or to rank sequences in terms of binding affinity (ranking).

To identify conditions in which machine learning with existing quantum annealing devices may be of use for studying a simplified biological problem, we report results obtained by solving a learning protocol with six different strategies: (i) an adiabatic quantum machine learning approach formulated in Refs. [4, 5] (DW), (ii) simulated annealing [48] (SA) (using the implementation given in Ref. [49]), (iii) simulated quantum annealing [50] (SQA), a classical algorithm that can represent the potential of a noiseless thermal quantum annealer, (iv) L_2 regularized multiple linear regression (MLR), (v) Lasso [51] and (vi) a scalable machine learning tool known as XGBoost (XGB) [52]. DW, SA and SQA are probabilistic approaches. SQA is a (classical) path integral Monte Carlo method that has performed very similarly to quantum annealing and captures some of its main advantages [53]. MLR is a deterministic method with a closed-form solution that returns the weights that best minimize the objective function (defined below). Lasso is a method for linear regression that uses an L_1 norm (see description of objective function below for more details). XGB uses boosted trees and has been applied to a variety of machine learning tasks in physics, natural language processing and ad-click prediction (e.g., Ref. [54]).

Given a transformed feature vector $\vec{\phi}_n$ that represents a DNA sequence, the goal of each method is to compute a predicted binding score $f(\vec{\phi}_n)$ that best matches the actual binding score. To carry out the task, an *objective function* must be optimized. The objective function consists of two parts: a training loss function and a regularization term that helps avoid overfitting. We may write the objective function as

$$Obj(\vec{w}) = R(\vec{w}) + \Omega(\vec{w}), \quad (2)$$

where R is the training loss, Ω is the regularization term, and \vec{w} is the set of feature weights to be determined by the six learning algorithms: DW, SA, SQA, MLR, Lasso and XGB. The

mean squared error was used as the loss function for all six methods; namely, $R(\vec{w}) = \sum_n (y_n - f_{\vec{w}}(\vec{\phi}_n))^2$, where y_n is the actual binding score of the n th sequence, and $f_{\vec{w}}(\vec{\phi}_n)$ is the predicted binding score. The regularization term was $\Omega(\vec{w}) = \lambda \|\vec{w}\|_1$ for DW, SA, SQA and Lasso, $\Omega(\vec{w}) = \lambda \|\vec{w}\|_2^2$ for MLR, and $\Omega(\vec{w}) = \gamma S + \frac{1}{2} \lambda \sum_{j=1}^S w_j^2$ for XGB, where the $\|\cdot\|_1$ norm is the number of 1's (Hamming weight), the $\|\cdot\|_2^2$ norm is the square of the Euclidean norm, and S is the number of leaves [52]. The calibration of the hyper-parameters λ , γ and S is discussed below. The loss function should be minimized and the regularization term generally controls model complexity by penalizing complicated models; the strength of the regularization was determined using a 100-fold Monte Carlo cross-validation. All six methods assume a linear model for the predicted binding affinity, i.e., $f_{\vec{w}}(\vec{\phi}_n) = \vec{w}^T \vec{\phi}_n = \sum_j w_j \phi_{n,j}$. DW, SA and SQA return binary weights and are probabilistic methods, that is, they return a distribution of weights with different energies [values of the Hamiltonian in Eq. (1)]. In order to utilize the distribution of weights returned, while not sacrificing the discrete nature of the QUBO approach, up to twenty of the best weights were averaged (see Supplementary Material, Sec. SID for a description of how excited-state solutions were included and Fig. S2 for an example of the decrease in the objective function).

Our computational procedure consisted of three main phases: (1) calibration of hyper-parameters, (2) training, and (3) testing (Fig. 2). About 10% of the data were held out for testing during the testing phase (“test data” or $\mathcal{D}^{\text{TEST}}$); this test data was not seen during calibration and training stages. Calibration and training were carried out using the remaining 90% of the data (“training data” or $\mathcal{D}^{\text{TRAIN}}$). Due to the discrete nature of the weights returned in the QUBO approach, as well as technological limitations of the DW2X device, calibration of hyper-parameter λ was carried out by repeatedly sampling a small number of sequences, about 2% and 10% of $\mathcal{D}^{\text{TRAIN}}$, corresponding to about 30 and 150 sequences, respectively. In particular, in the calibration phase we determined the hyper-parameters by using 100-fold Monte Carlo (or split and shuffle) cross-validation with training splits of 2% and 10% of the training data, varying λ from 2^{-3} to 2^6 . Monte Carlo cross-validation was used so that hyper-parameters would be tuned on a similar number of sequences as used in the training phase (in contrast, n -fold cross-validation trains on $\frac{n-1}{n} \times 100\%$ of the data). The same calibration procedure was applied to tune λ for SA, SQA, MLR and Lasso: the resulting values of λ are listed in the Supplementary Material, Tables S1 and S2. In order to demonstrate good performance for XGB, γ , S , and several additional parameters needed to be tuned (see Methods). In the training phase we used a bagging (bootstrap aggregating) procedure [55], randomly sampling with replacement 2% and 10% of the training data, namely about 30 and 150 sequences. Each subset of about 30 or 150 sequences formed a training “instance”, and the mapping of a subset of data to the h_i and J_{ij} seen by DW and SA is given in the Methods. Each learning approach (DW, SA, SQA, MLR, Lasso and XGB) was trained on the same set of instances. To collect statistics, 50 instances were randomly selected with replacement, for each training size. In the testing phase, the predictive power was assessed in terms of classification performance (the mean area under the precision-recall curve or AUPRC) and ranking performance (the median Kendall’s τ) on the test data unseen during calibration and training phases. AUPRC is a measure of classification performance that may help discern between similar algorithms when there is a high degree of class

imbalance; i.e., when the data set contains many more false labels than true labels [56]. Kendall's τ is a rank correlation coefficient that counts all mismatches equally [57]. Additional methodological details are given in Methods and in the Supplementary Material, Sec. SI.

Performance on gcPBM Data

To quantify the relative performance of the algorithms in capturing DNA-protein binding preferences, we first present results for high-quality gcPBM [40] data of three TFs from the basic helix-loop-helix (bHLH) family: the Mad1/Max heterodimer ('Mad'), the Max homodimer ('Max'), and the c-Myc/Max heterodimer ('Myc') [44]. bHLH proteins typically recognize and bind as dimers to the enhancer box (E-box), which is of the form CANNTG, where N denotes any of the 4 nucleotides (A, C, G, or T). Mad, Max, and Myc are part of a gene network that controls transcription in cells; a mutation of Myc has been associated with many forms of cancer [58]. For the work here, these three datasets were modified to consist of about 1600 sequences of 10 base pairs (bp) in length with the E-box located at the central 6 bp.

In Fig. 3 we present the AUPRC and Kendall's τ obtained with the different algorithms when training with about 30 (2%) and 150 (10%) sequences. To compute the AUPRC, a threshold of the data was introduced: for a threshold at the p th percentile of the data, $p\%$ of the total number of sequences have binding affinities below the threshold and were set as negatives ("false"), and the $(1 - p)\%$ of the sequences that have binding affinities above the threshold were set as positive ("true"); see Supplementary Material, Sec. SID for a more detailed explanation of the procedure to threshold the data and to generate and calculate the AUPRC. During the calibration phase, we tuned hyper-parameters with a single threshold at the 80th percentile of the data, and during the testing phase we evaluated performance between the 70th and the 99th percentiles of the data. Kendall's τ was evaluated between the predicted and measured binding affinity. A higher AUPRC indicates a better ability to correctly classify sequences that would be strongly bound by a TF, and a higher τ indicates a better ability to accurately rank the binding affinities for different sequences.

For the AUPRC, when training on instances with 2% of the data (left column in Fig. 3a), DW, SA and SQA perform very similarly, with DW slightly outperforming SA on the Myc data, and are somewhat better than MLR at the 70th and 80th percentiles. MLR tends to do better at the higher thresholds: this behavior could be affected by the fact that, during the calibration phase, we selected the λ that gave the best performance at the 80th percentile. Lasso, which uses the same L_1 norm as DW, SA, and SQA, performs better than XGB but worse than the other methods. XGB, which has been successfully applied to a growing number of learning tasks, does poorly with small training sizes. When training with 10% of the data (right column in Fig. 3a), the trends of relative classification performance are quite different. XGB and MLR perform very similarly, though XGB does slightly better for the Max dataset. DW tends to perform better than SA and SQA, especially at higher thresholds. DW's mean AUPRC is normally worse than MLR and XGB's, though there is overlap between the error bars. SA and SQA generally perform worse than the other methods, but not conspicuously so. A more thorough analysis of DW's classification performance in

comparison to SA and SQA with the same problem parameters is reported in the Supplementary Material, Figs. S3–S6 and related text in Sec. SII A. Lasso's performance is in general comparable to DW, SA, and SQA and generally seems to perform the worst with 10% of training data.

For Kendall's τ (Fig. 3b), Lasso and XGB's performance are the least favorable when training with 2% of the data. SQA generally performs the best over the three TFs, though MLR's median τ is marginally greater than SQA's for Mad. SA's performance is very close to SQA's and DW's performance is slightly worse than the other two annealing schemes, though generally better than the typical machine learning algorithms. With 10% of the data, DW performs the worst; SA and SQA perform very similarly, with SQA being slightly better on two of the three datasets; MLR and Lasso perform very similarly, though MLR looks slightly better; and XGB performs the best.

The fact that for Mad and Max with 2% of the training data there is very little variation in Kendall's τ for SA and SQA (and to a lesser extent, DW), is a consequence of the choice of hyper-parameters. The specific values of the hyper-parameters that gave optimal value of Kendall's τ during the calibration phase are shown in Supplementary Material, Table S2, but we note here that the value of λ is quite high. λ controls the model complexity and is closely related to the bias-variance tradeoff, which states that it is impossible to simultaneously minimize errors from both bias and variance. A large value of λ introduces a large bias [59]; consequently, for the cases where there is no or little variance, SA and SQA are essentially extracting the same pattern from all the training data. For the ranking tasks shown here with training on about 30 sequences, this gives the best performance for SA and SQA. It may be unsurprising, however, that a large value of λ be appropriate for small datasets; over-fitting may be a greater concern with smaller amounts of data.

The results presented in Fig. 3 suggest a precise case where current quantum technology may offer slight performance advantages relative to classical computational approaches; that is, when there is only a small amount of experimental training data available (about 30 sequences in our specific cases). In both classification and ranking tasks, DW performs comparably to SA and SQA and better than Lasso and XGB. MLR performs comparably with the annealing methods, but its error bars are much larger, indicating that its performance is less stable and more dependent on the training data. Moreover, the similarity between DW and SQA suggests that for small training sizes DW is functioning very nearly like a noiseless quantum annealer as captured by quantum Monte Carlo simulations. On a larger size of the training data DW's performance decreases relative to the classical approaches for all three TFs, though results are still competitive. The decrease in the performance of all annealing methods (DW, SA, and SQA) seems to indicate a limitation on using methods with discrete weights, which enforce simpler models. Such models may be more advantageous with a small number of training samples because they prevent overfitting. However, with larger amounts of training data, a simpler model may not adequately learn the variation within the data and hence suffer worse performance. Nevertheless, the fact that Lasso uses the same L_1 norm as the annealing methods (i.e., DW, SA and SQA), yet does not perform as well, indicates an advantage of such annealing

methods when training with a small number of sequences. This is consistent with the finding reported in Ref. [18].

Weight Logos from Feature Weights

Since the one-hot encoding was used with a linear model to represent DNA sequence, the weights returned by DW, SA, and MLR reflect the relative importance of the corresponding nucleotide at a position in the sequence for the binding score. The magnitude of these feature weights for DW, SA, and MLR can be visualized as a “weight logo” and are presented in Fig. 4 for the Mad, Max, and Myc gcPBM datasets. XGB, which finds an ensemble of trees, does not assign weights to individual nucleotides and hence does not easily lend itself to visualization. Similar plots for SQA and Lasso are shown in the Supplementary Material Sec. SIIB and Fig. S7. The weight logos show the contribution of nucleotides at particular positions to the strength of binding. The contribution of a nucleotide at a particular position in the sequence is represented by its height; nucleotides with the smallest weights are at the bottom and those with the largest weights are at the top. These weight logos in Fig. 4 were obtained by averaging the weights from the 50 training instances of the same number of sequences with the AUPRC as the objective. In other words, the logo represents the average of the weights that give the AUPRCs shown in Fig. 3a. DW, SA, and MLR all perform very similarly and give weight logos that are in good agreement with the expected consensus sequence, CANNTG. This demonstrates that all methods are able to capture biologically relevant information.

Performance on HT-SELEX Data

HT-SELEX [41, 43] is a method for investigating the relative binding affinity of a TF for a particular sequence of DNA, an in vitro technique complementary to PBM. We present results for the Max homodimer and TCF4, another member of the bHLH family with consensus sequence CANNTG, using data from HT-SELEX experiments [47]. The Max dataset consisted of 3200 sequences of 12 bp in length, and the TCF4 dataset was modified to contain 1800 sequences of 12 bp in length.

The procedure for splitting each dataset into test and training data was similar to that described earlier for the gcPBM datasets (see Fig. 2). There was no overlap between training and testing data. The quantitative results for classification and ranking performance of the six different machine learning approaches are summarized in Figs. 5a and 5b, and the weight logos for DW, SA, and MLR in Fig. 5c. As with the gcPBM data, when training with about 30 sequences (1% of training data for Max and 2% of the training data for TCF4), DW, SA and SQA exhibit the best performance on the test dataset. MLR matches the annealing protocols with a threshold at the 70th and 80th percentile of the data, but does worse at the higher percentiles of the data (left column in Figs. 5a and 5b). Lasso and XGB have the poorest performance. When training with about 150 (5% of training data for Max and 10% of the training data for TCF4) sequences, XGB performs very well, as on the gcPBM datasets with more training data, and MLR does well on the Max dataset but rather poorly on the TCF4 dataset; Lasso is comparable to MLR. DW’s performance is worse than the best performing method (XGB), but comparable to the other methods (right column in Fig.

5a and 5b). XGB's performance on the TCF4 dataset is much better than the other methods, except when thresholding at the 99th percentile.

In terms of Kendall's τ (Fig. 5a and 5b, bottom), all methods have similar performance when training with about 30 sequences, with the exception of XGB which does not do as well on the Max dataset. When training with about 150 sequences, XGB gives the best ranking performance, as it did with the gcPBM data, and the other methods all perform similarly. Finally, the weight logos in Fig. 5c indicate that DW, SA, and MLR capture patterns in the data that give good agreement with the expected consensus sequence. The weight logos for the Max and TCF4 HT-SELEX datasets from SQA and Lasso are reported in Supplementary Material Fig. S8.

DISCUSSION

In this work we have explored the possibility of using a machine learning algorithm based on quantum annealing to solve a simplified but actual biological problem, the classification and ranking of TF-DNA binding events. This is the first application of quantum annealing to real biological data.

We have shown that DW performs comparably or slightly better than classical counterparts for classification when the training size is small, and competitively for ranking tasks. This trend is consistent with results on older sets of gcPBM and HT-SELEX data for various TFs, which are reported in Supplementary Material, Sec. SIII, Figs. S9–S14. Moreover, these results are consistent with a similar approach for the Higgs particle classification problem [18], where DW and SA both outperformed XGB with small training sizes, with a slight occasional advantage for DW over SA. This robustness across completely different application domains suggests that these findings represent real present-day advantages of annealing approaches over traditional machine learning in the setting of small-size training data. In areas of research where datasets with a small number of relevant samples may be more common, a QUBO approach such as quantum annealing realized via DW may be the algorithm of choice. On the other hand, when data is plentiful, some of the other state-of-the-art classical algorithms may be a better choice.

We have also demonstrated that the feature weights obtained by DW reflect biological data; the weight logos for the TF-DNA binding data from gcPBM and HT-SELEX are consistent with the consensus binding site. This gives some confidence that quantum annealing is learning relevant biological patterns from the data. Yet, the approach is not without limitations. One limitation comes from the use of a single-nucleotide model to encode the DNA binding sites. In fact, we implicitly used a simple model that assumes independence between positions in the sequence. This is not always a valid approximation; higher-order “k-mer” features or other “shape” features that account for interdependencies between nucleotide positions may enhance model precision [44, 45, 47, 60, 61]. We are limited to this simple model because of major technological constraints on the number of available qubits, which limits the number of features that can be used and thus the length of sequences that can be examined. The DW2X processor used for this study has 1098 functional qubits, but because of a sparse connectivity between qubits, only 40 or so features can actually be

implemented on the device and in our study (see Supplementary Material, Sec. SIA for more details). Another serious limitation is the use of discrete weights. Discrete weights seem to be advantageous with a small number of training samples, as they enforce simpler models and are less prone to overfitting. However, as the amount of training data increases, these simpler models do not fare as well as some of the classical methods, which allow for greater numerical precision in the weights.

Despite these limitations, it is encouraging to see competitive performance for the simplified problem we have studied here. Although the performance advantage from annealing-type optimizers makes it difficult to solely attribute the performance to quantumness, this work may inspire future investigations into the power of quantum annealing devices. As quantum technology continues to develop and advance, it is possible that some of the practical limitations will be addressed and the range of problems that can be explored will be expanded.

METHODS

QUBO mapping of TF-DNA binding problem

After processing the experimental datasets of N sequences of fixed length L and a measure of the binding affinity, we obtained the restricted datasets to which we applied six different machine learning strategies. Datasets were formulated as $\{(\vec{\phi}_n, y_n)\}_{n=1}^N$, where $\vec{\phi}_n \equiv (\phi_{n,1}, \dots, \phi_{n,4L})^\top$ is the transformed feature vector, and y_n is the binding affinity. Solving for the simplest model is equivalent to finding a vector of binary weights $\vec{w} = (w_1, \dots, w_{4L})$, where $w_i \in \{0, 1\}$, such that the quantity

$$\delta = \sum_{n=1}^N \left(y_n - \vec{w}^\top \vec{\phi}_n \right)^2 \quad (3)$$

is minimized. The problem can then be specified as finding a \vec{w}_{opt} such that

$$\vec{w}_{\text{opt}} = \arg \min_{\vec{w}} \sum_{n=1}^N \left(y_n - \vec{w}^\top \vec{\phi}_n \right)^2 + \lambda \|\vec{w}\|_1, \quad (4)$$

where λ is a regularization (penalty) term included to prevent overfitting and $\|\vec{w}\|_1 = \sum_m w_m$ is the number of non-zero weights. To represent the above as an Ising problem, note that we can rewrite Eq. (4) as follows:

$$\begin{aligned} \vec{w}_{\text{opt}} &= \arg \min_{\vec{w}} \sum_{n=1}^N \left(y_n - \vec{w}^\top \vec{\phi}_n \right)^2 + \lambda \sum_{m=1}^{4L} w_m \\ &= \arg \min_{\vec{w}} \sum_n \left(y_n^2 - 2y_n \vec{w}^\top \vec{\phi}_n + \vec{w}^\top \vec{\phi}_n \vec{\phi}_n^\top \vec{w} \right) + \lambda \sum_{m=1}^{4L} w_m = \arg \min_{\vec{w}} \vec{w}^\top Q \vec{w} + \vec{w}^\top k, \end{aligned} \quad (5)$$

where

$$Q = \sum_n \vec{\phi}_n \vec{\phi}_n^\top; Q_{i,j} = \sum_n \phi_{n,i} \phi_{n,j} \quad (6)$$

$$k = \lambda \mathbf{1} - 2 \sum_n y_n \vec{\phi}_n; k_i = \lambda - 2 \sum_n y_n \phi_{n,i}.$$

Constants that do not affect the optimization are dropped in the latter step. This procedure demonstrates that the problem of TF-DNA binding can be formulated as a QUBO problem, which in turn can easily be transformed into an Ising Hamiltonian of the form in Eq. (1) and passed to D-Wave. The data normalization procedure is described in Supplementary Material, Sec. SI C.

Technical details of algorithms

In order to solve practical problems of interest on DW, an embedding procedure must be used (see Supplementary Material, Sec. I A). Some additional preprocessing was also performed for DW and SA to ensure that all response values were feasible (see Supplementary Material, Sec. SI C). DW, SA, SQA, MLR, Lasso and XGB were run on the same set of instances for assessment of the quantum annealer on the chosen problem. The experimental quantum processor, DW2X, was designed and built by D-Wave Systems, Inc. For each instance a total of 10000 anneals (“runs”) were collected from the processor, run with an annealing time of 20 μ s. SA and SQA are classical analogues of quantum annealing that perform annealing on a classical and path integral Monte Carlo simulation of the Ising spin glass, respectively. SA and SQA were run with 10000 sweeps (each sweep is an update of all spins) per repetition (or “anneals”) with an initial inverse temperature of 0.1 and a final inverse temperature of 3, for a total of 10000 repetitions. The SA code was adapted from Ref. [49], and an in-house version of SQA was used. MLR is a widely used technique to minimize the loss function shown in Eq. (4), with the convex penalty term $\lambda \|\vec{w}\|_2^2$ instead of the linear penalty term. Lasso has the linear penalty term [51]; XGB uses boosted trees [52]. The weights \vec{w} returned by MLR, Lasso and XGB are real-valued, whereas the weights returned by DW, SA and SQA (which solve a QUBO/Ising problem) are binary. In addition, DW, SA and SQA are probabilistic, meaning that a distribution of weights with different energies [the value of H_P in Eq. (1)] are returned. Up to 20 of the lowest energy weights were included for both DW, SA and SQA (see Supplementary Material, Sec. SID for more details). The lower the energy, the better the particular solution is at minimizing Eq. (4). In contrast, MLR, Lasso and XGB are deterministic and return a single solution.

In the calibration phase, only one hyper-parameter, λ was tuned for DW, SA, SQA, MLR and Lasso. All five methods were tuned separately for both classification and ranking tasks, resulting in different optimal λ for each method (see Supplementary Material Tables S2 and S3 for final values of λ). With an older dataset we varied both the number of sweeps for SA and the value of λ but results were not significantly different; hence, here we only vary λ for

SA. SA also has various other parameters that are related to the algorithm itself, including number of runs, initial and final temperature, and the cooling schedule, all of which affect the optimization performance. These parameters were not tuned. Similar additional parameters for DW, including annealing time and number of runs, were also not tuned. XGB's performance depends on several hyper-parameters, and more careful tuning was necessary in order to give competitive performance. XGB parameters [52] that were considered include γ , the max depth, and min child weight (all of which control model complexity), subsample, colsample bytree, (which add randomness to make training robust to noise), as well as learning rate, etc. Rather than doing a full grid search over all these parameters, parameters were tuned sequentially; i.e., one value of η was fixed, then the best value of max depth and min child weight were found. The optimal γ for those values was then found; and finally subsample and colsample bytree tuned. η was then varied and the process repeated. η was varied from 0.05 to 0.3, max depth from 3 to 20, min child weight from 1 to 20, γ from 0 to 1, and subsample and colsample bytree both from 0.6 to 1.

In the testing phase, we evaluated performance based on two metrics: the AUPRC for classification performance and Kendall's τ for ranking performance. For the AUPRC, we reported mean values with standard deviations as error bars, whereas for Kendall's τ the median value was presented.

Data processing and availability

Original probes for the gcPBM [40] data contained 16000–18000 sequences of 36 bp in length with the fluorescence intensity as a measure of binding affinity. The same data is used in [44] and may be downloaded from GEO (<https://www.ncbi.nlm.nih.gov/geo/>) under accession number GSE59845. Because of current limitations of the architecture of the DW device that limit the number of features that may be used, the data was truncated to the central 10 bp. For each sequence of 10 bp, we calculated its average gcPBM signal. In other words, all sequences in the datasets were unique. The final Mad, Max, and Myc datasets consisted of 1655, 1642, and 1584 sequences, respectively, of length 10 bp, and the logarithm base 2 with fluorescence intensities was used. The HT-SELEX data came from mammalian TFs [42] that was re-sequenced with on average 10-fold increase in sequencing depth [47]. The sequencing data is available at the European Nucleotide Archive (ENA - <https://www.ebi.ac.uk/ena>; study identifier PRJEB14744) and was pre-processed following the protocol in [47]. After this first step of pre-processing, the Max and TCF4 datasets consisted of 3209 and 15556 sequences of length 12 bp and 14 bp, respectively. The Max dataset did not require further truncation, but one bp on the left and right flanks were trimmed for the TCF4 dataset, giving a modified dataset of 1826 sequences of length 12 bp. As with the gcPBM data, the average relative affinity was averaged for each truncated sequence.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Lin Yang for assistance with pre-processing the gcPBM and HT-SELEX data. The authors would also like to thank Tameem Albash for providing the SQA code. This work was supported by the USC Women in Science and Engineering Program (to R.D.F), the National Institutes of Health grants R01GM106056 and U01GM103804 (to R.R.), ARO grant number W911NF-12-1-0523 and NSF grant number INSPIRE-1551064 (to D.L.). R.R. is an Alfred P. Sloan Research Fellow. D.L. is a Guggenheim Foundation Fellow and a Caltech Moore Scholar in Physics. This research is based upon work (partially) supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. Army Research Office contract W911NF-17-C-0050. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

References

1. Shor PW. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* 1997; 26:1484.
2. Feynman RP. Quantum mechanical computers. *Optics News.* 1985; 11:11–20.
3. Lloyd S. Universal quantum simulators. *Science.* 1996; 273:1073–1078. [PubMed: 8688088]
4. Neven H, Denchev VS, Rose G, Mcready WG. Training a binary classifier with the quantum adiabatic algorithm. *arXiv:0811.0416.* 2008
5. Pudenz KL, Lidar DA. Quantum adiabatic machine learning. *Quantum Information Processing.* 2013; 12:2027–2070.
6. Lloyd S, Mohseni M, Rebentrost P. Quantum algorithms for supervised and unsupervised machine learning. *arXiv:1307.0411.* 2013
7. Wittek, P. *Quantum Machine Learning: What Quantum Computing Means to Data Mining.* Elsevier, AP; 2014.
8. Lloyd S, Mohseni M, Rebentrost P. Quantum principal component analysis. *Nat. Phys.* 2014; 10:631–633.
9. Rebentrost P, Mohseni M, Lloyd S. Quantum support vector machine for big data classification. *Phys. Rev. Lett.* 2014; 113:130503–. [PubMed: 25302877]
10. Wiebe N, Kapoor A, Svore KM. Quantum deep learning. *arXiv:1412.3489.* 2014
11. Schuld M, Sinayskiy I, Petruccione F. An introduction to quantum machine learning. *Contemporary Physics.* 2015; 56:172–185.
12. Wiebe N, Kapoor A, Svore K. Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning. *Quantum Information & Computation.* 2015; 15:0318–0358.
13. Aaronson S. Read the fine print. *Nat. Phys.* 2015; 11:291–293.
14. Adachi SH, Henderson MP. Application of quantum annealing to training of deep neural networks. *arXiv:1510.06356.* 2015
15. Amin MH, Andriyash E, Rolfe J, Kulchysky B, Melko R. Quantum boltzmann machine. *arXiv: 1601.02036.* 2016
16. Biamonte J, et al. Quantum machine learning. *arXiv:1611.09347.* 2016
17. Benedetti M, Realpe-Gómez J, Biswas R, Perdomo-Ortiz A. Quantum-assisted learning of graphical models with arbitrary pairwise connectivity. *arXiv:1609.02542.* 2016
18. Mott A, Job J, Vlimant J-R, Lidar D, Spiropulu M. Solving a higgs optimization problem with quantum annealing for machine learning. *Nature.* 2017; 550:375 EP –. [PubMed: 29052620]
19. Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. *Phys. Rev. E.* 1998; 58:5355.
20. Farhi E, et al. A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem. *Science.* 2001; 292:472–475. [PubMed: 11313487]
21. Albash T, Lidar DA. Adiabatic quantum computing. *arXiv:1611.04471.* 2016
22. Harris R, et al. Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor. *Phys. Rev. B.* 2010; 82:024511.

23. Johnson MW, et al. Quantum annealing with manufactured spins. *Nature*. 2011; 473:194–198. [PubMed: 21562559]
24. Bunyk PI, et al. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Transactions on Applied Superconductivity*. Aug.2014 24:1–10.
25. Kato T. On the adiabatic theorem of quantum mechanics. *J. Phys. Soc. Jap.* 1950; 5:435.
26. Farhi E, Goldstone J, Gutmann S, Sipser M. Quantum Computation by Adiabatic Evolution. arXiv:quant-ph/0001106. 2000
27. Jansen S, Ruskai M-B, Seiler R. Bounds for the adiabatic approximation with applications to quantum computation. *J. Math. Phys.* 2007; 48:102111.
28. Childs AM, Farhi E, Preskill J. Robustness of adiabatic quantum computation. *Phys. Rev. A*. 2001; 65:012322.
29. Amin MHS, Averin DV, Nesteroff JA. Decoherence in adiabatic quantum computation. *Phys. Rev. A*. 2009; 79:022107.
30. Albash T, Lidar DA. Decoherence in adiabatic quantum computation. *Phys. Rev. A*. 2015; 91:062320–.
31. Barahona F. On the computational complexity of Ising spin glass models. *J. Phys. A: Math. Gen.* 1982; 15:3241.
32. Rønnow TF, et al. Defining and detecting quantum speedup. *Science*. 2014; 345:420–424. [PubMed: 25061205]
33. Katzgraber HG, Hamze F, Andrist RS. Glassy chimeras could be blind to quantum speedup: Designing better benchmarks for quantum annealing machines. *Phys. Rev. X*. 2014; 4:021008–.
34. Venturelli D, et al. Quantum optimization of fully connected spin glasses. *Phys. Rev. X*. 2015; 5:031040–.
35. Hen I, et al. Probing for quantum speedup in spin-glass problems with planted solutions. *Phys. Rev. A*. 2015; 92:042325–.
36. Amin M. Searching for quantum speedup in quasistatic quantum annealers. arXiv:1503.04216. 2015
37. Slattery M, et al. Absence of a simple code: how transcription factors read the genome. *Trends in Biochemical Sciences*. 2014; 39:381–399. [PubMed: 25129887]
38. Shlyueva D, Stampfel G, Stark A. Transcriptional enhancers: from properties to genome-wide predictions. *Nat Rev Genet*. 2014; 15:272–286. [PubMed: 24614317]
39. Stormo G, Zhao Y. Determining the specificity of dna-protein interactions. *Nature Reviews Genetics*. 2010; 11:751–760.
40. Gordân R, et al. Genomic regions flanking e-box binding sites influence DNA binding specificity of bHLH transcription factors through DNA shape. *Cell Reports*. 2013; 3:1093–1104. [PubMed: 23562153]
41. Jolma A, et al. Multiplexed massively parallel selex for characterization of human transcription factor binding specificities. *Genome Research*. 2010; 20:861–873. [PubMed: 20378718]
42. Jolma A, et al. Dna-binding specificities of human transcription factors. *Cell*. 2013; 152:327–339. [PubMed: 23332764]
43. Liu J, Stormo GD. Combining selex with quantitative assays to rapidly obtain accurate models of protein-dna interactions. *Nucleic Acids Res*. 2005; 33:e141. [PubMed: 16186128]
44. Zhou T, et al. Quantitative modeling of transcription factor binding specificities using DNA shape. *Proceedings of the National Academy of Sciences*. 2015; 112:4654–4659.
45. Yang L, et al. TFBSshape: a motif database for DNA shape features of transcription factor binding sites. *Nucleic Acids Research*. 2014; 42:D148–D155. [PubMed: 24214955]
46. Abe N, et al. Deconvolving the recognition of DNA shape from sequence. *Cell*. 2015; 161:307–318. [PubMed: 25843630]
47. Yang L, et al. Transcription factor family-specific DNA shape readout revealed by quantitative specificity models. *Molecular Systems Biology*. 2017; 13
48. Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science*. 1983; 220:671–680. [PubMed: 17813860]

49. Isakov S, Zintchenko I, Rønnow T, Troyer M. Optimized simulated annealing for ising spin glasses. arXiv:1401.1084. 2014
50. Santoro GE, Marto ́ak R, Tosatti E, Car R. Theory of quantum annealing of an Ising spin glass. *Science*. 2002; 295:2427–2430. [PubMed: 11923532]
51. Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1996:267–288.
52. Chen, T., Guestrin, C. Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining KDD '16. New York, NY, USA: ACM; 2016. Xgboost: A scalable tree boosting system; p. 785-794.
53. Crosson E, Harrow A. Simulated quantum annealing can be exponentially faster than classical simulated annealing. arXiv preprint arXiv:1601.03030. 2016
54. Chen, T., He, T. Higgs boson discovery with boosted trees. In; Proceedings of the 2014 International Conference on High-Energy Physics and Machine Learning - Volume 42, HEPML'14. 2014. p. 69-80.JMLR.org
55. Breiman L. Bagging predictors. *Machine learning*. 1996; 24:123–140.
56. Davis, J., Goadrich, M. Proceedings of the 23rd International Conference on Machine Learning, ICML '06. New York, NY, USA: ACM; 2006. The relationship between precision-recall and roc curves; p. 233-240.
57. Kendall M. A new measure of rank correlation. *Biometrika*. 1938; 30:81–93.
58. Grandori C, Cowley S, James L, Eisenman R. The Myc/Max/Mad network and the transcriptional control of cell behavior. *Annual Review of Cell and Developmental Biology*. 2000; 16:653–699.
59. Cucker F, Smale S. Best choices for regularization parameters in learning theory: on the bias-variance problem. *Foundations of Computational Mathematics*. 2002; 2:413–428.
60. Mordelet F, Horton J, Hartemink AJ, Engelhardt BE, Gordán R. Stability selection for regression-based models of transcription factor–dna binding specificity. *Bioinformatics*. 2013; 29:i117–i125. [PubMed: 23812975]
61. Zhou T, et al. DNASHape: a method for the high-throughput prediction of DNA structural features on a genomic scale. *Nucleic Acids Research*. 2013; 41:W56–W62. [PubMed: 23703209]

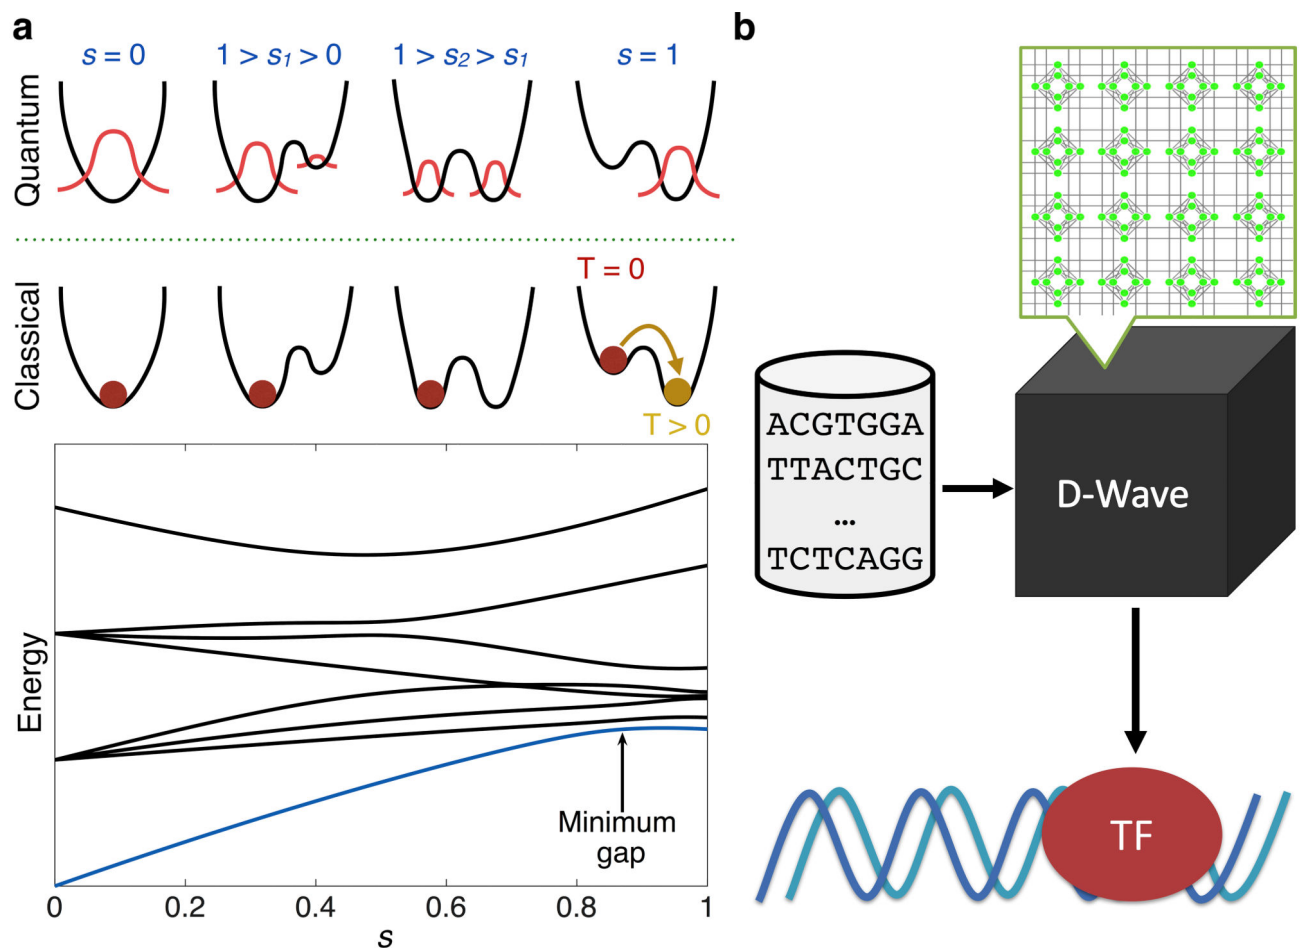


FIG. 1. Illustration of the principles and purpose of this work. **(a)** Top: Quantum versus classical adiabatic annealing processes. With quantum annealing there is the possibility for the system state (red) to tunnel through a changing barrier (black) and arrive at the ground state; for classical annealing, the system must rely on thermal fluctuations (temperature $T > 0$) to overcome any energy barriers. Bottom: Typical spectrum of instantaneous energy eigenvalues during adiabatic quantum optimization. The ground state energy at $s = 0$ has a significant gap to the next energy level. The speed at which the optimization can take place depends on the size of the minimum gap. **(b)** Using simplified datasets of a small number of sequences derived from actual binding affinity experiments, we use D-Wave, a commercially available quantum annealer, to classify and rank binding affinity preferences of a TF to DNA.

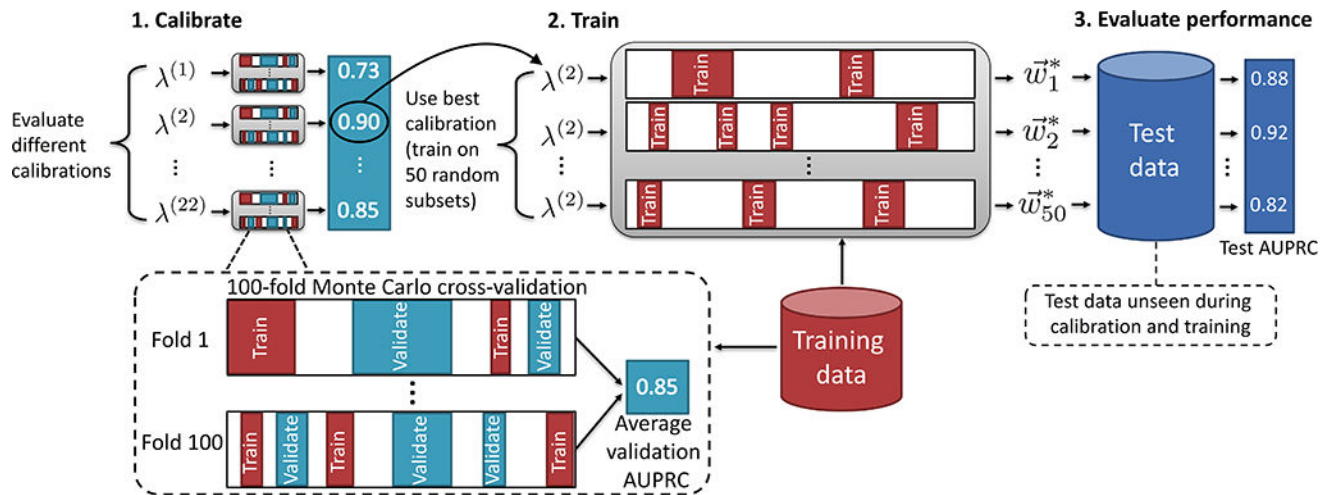
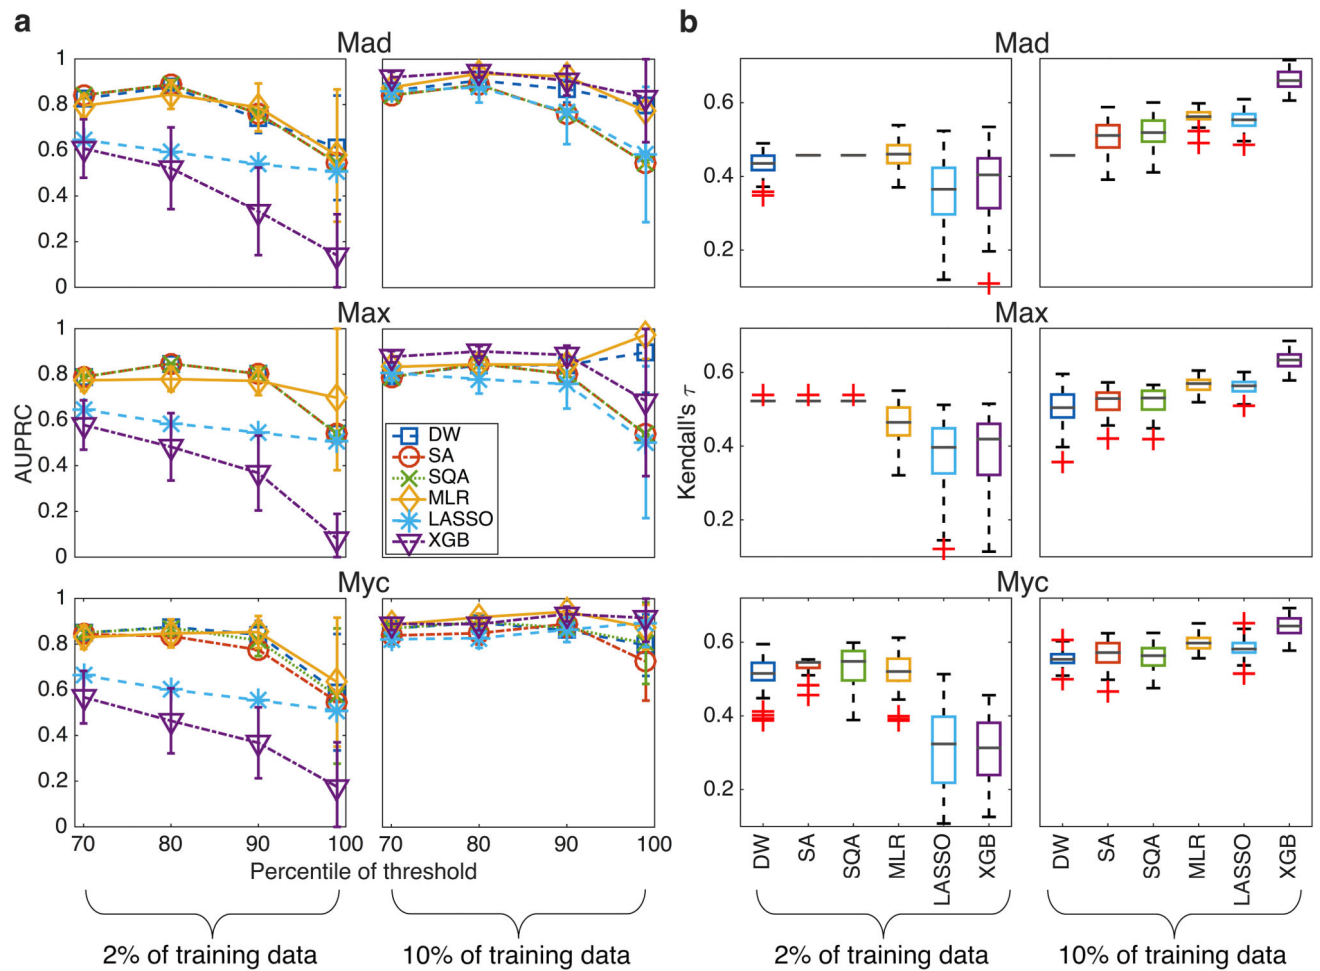


FIG. 2.

Schematic overview of the data handling procedure. About 10% of the data were held out for testing ($\mathcal{D}^{\text{TEST}}$) and was unseen during the calibration and training phases. The remaining 90% of the data were used as calibration and training ($\mathcal{D}^{\text{TRAIN}}$). For DW, SA, and MLR, in the calibration step, hyper-parameter λ was tuned using a 100-fold Monte Carlo cross-validation. For XGB, several other hyper-parameters were tuned (see Methods for a list). During the training step, a procedure similar to bagging (bootstrap aggregating) was used by randomly sampling a 2% and 10% replacement of the data 50 times to give 50 training instances. In the testing step, the area under the precision-recall curve (AUPRC) of the best performing weights for each of the 50 training instances was evaluated on $\mathcal{D}^{\text{TEST}}$, which was unseen during training and calibration, to evaluate generalization of classification performance. The calibration, training, and testing procedure was identical for ranking tasks, with the exception that Kendall's τ was used as the metric of performance instead of the AUPRC.

**FIG. 3.**

Quantitative performance on held-out experimental test dataset of two different types of tasks for three high quality gcPBM datasets. **(a)** The mean AUPRC for Mad, Max, and Myc plotted versus threshold at certain threshold percentiles of the data, when training with 2% of the data (left) and 10% of the data (right). In both cases 50 instances were randomly selected for training and performance of the 50 trained weights is evaluated on the same held-out test set. Error bars are the standard deviation. **(b)** Boxplot of Kendall's τ on held-out test dataset. Red '+' indicate outliers, gray line represents the median. The bottom and top edges of the box represent the 25th and 75th percentiles, respectively.

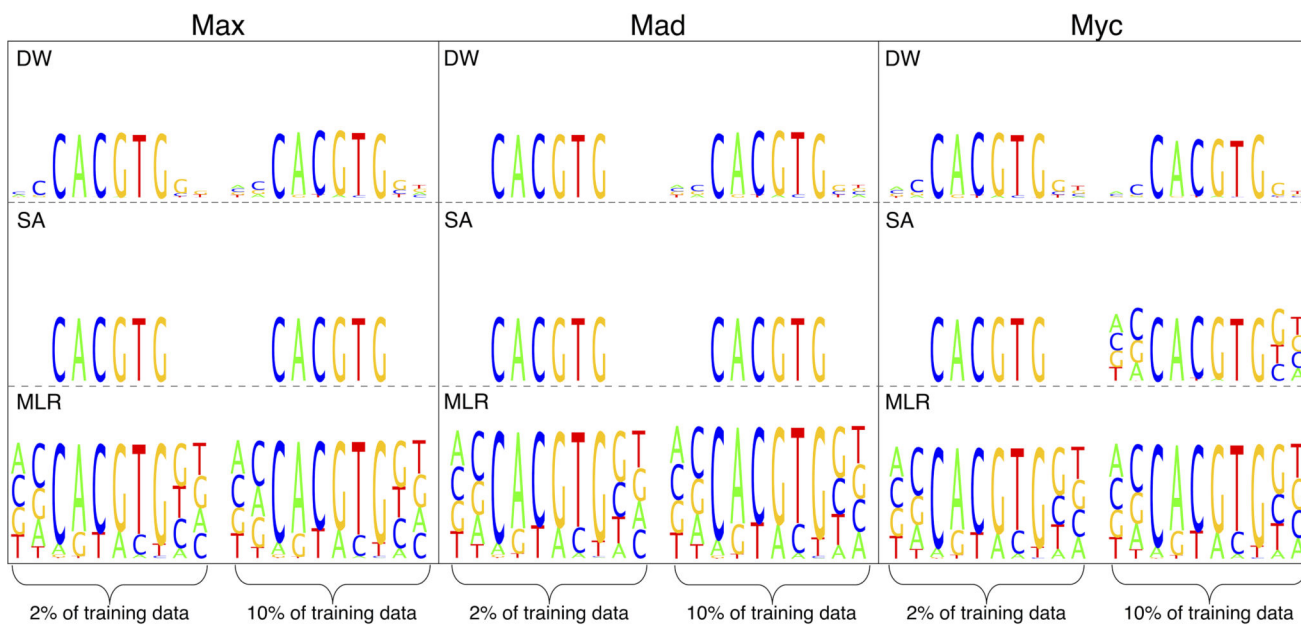
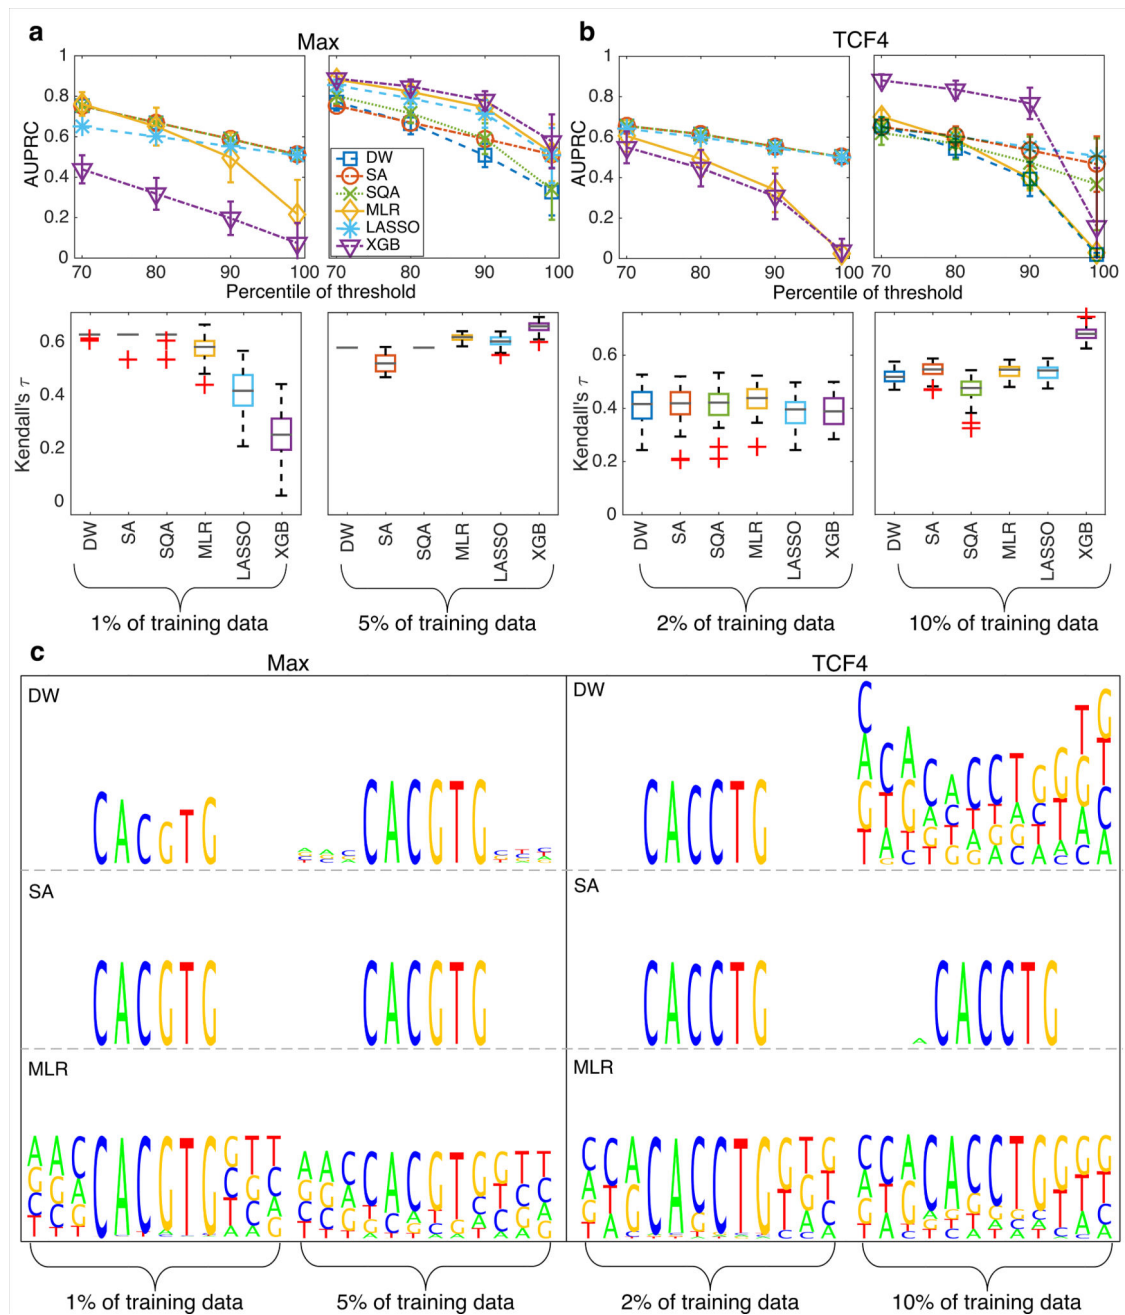


FIG. 4. Comparison of feature weights visualized as weight logos for DW, SA, and MLR. Weights represent the relative importance of a nucleotide at each position for the binding affinity. These weight logos were obtained using the Mad, Max, and Myc gcPBM datasets when training with the AUPRC as the objective.

**FIG. 5.**

Summary of results for HT-SELEX data. **(a)** Comparison of the AUPRC (top) and Kendall's τ (bottom) when training with 1% of the data (left) and 5% of the data (right) for Max. Error bars are standard deviation over 50 instances. **(b)** Comparison of the AUPRC (top) and Kendall's τ (bottom) when training with 2% of the data (left) and 5% of the data (right) for TCF4. **(c)** Weight logos for Max and TCF4 from training with the AUPRC as the objective.