

CONVEX RELAXATIONS OF CONVOLUTIONAL NEURAL NETS

Burak Bartan Mert Pilanci

Department of Electrical Engineering
Stanford University

ABSTRACT

We propose convex relaxations for convolutional neural nets with one hidden layer where the output weights are fixed. For convex activation functions such as rectified linear units, the relaxations are convex second order cone programs which can be solved very efficiently. We prove that the relaxation recovers the global minimum under a planted model assumption, given sufficiently many training samples from a Gaussian distribution. We also identify a phase transition phenomenon in recovering the global minimum for the relaxation.

Index Terms— Convolutional neural networks, convex relaxations, linear programming, deep learning

1. INTRODUCTION

Convolutional neural networks (CNNs) have been extremely successful across many domains in machine learning and computer vision [15, 16]. However, analyzing the behavior of non-convex optimization methods used to train CNNs remains a challenge. In this paper, we propose finite dimensional convex relaxations for convolutional neural nets which are composed of a single hidden layer with convex activation functions. We prove that for the ReLU activation function, a randomized perturbation of the convex relaxation recovers the global optimum with probability approaching $\frac{1}{2}$, which can be amplified to $1 - (\frac{1}{2})^r$ using r independent trials of random perturbations. We illustrate a phase transition phenomenon in the probability of recovering the global optimum. Finally, we consider learning filters for a regression task on the MNIST dataset of handwritten digits [17].

1.1. Related Work and Contributions

In recent years, there has been an increasing amount of literature on providing theoretical results for neural networks. A considerable amount of work in this area focused on the case where a convolutional network with a single hidden layer is trained using gradient descent. For instance, [4] shows that gradient descent achieves the global minimum (with high probability) for convolutional networks with one layer and no overlap when the distribution of the input is Gaussian. [4]

also proves that the problem of learning this network is NP-complete. Zhang et. al propose a convex optimization approach based on a low-rank relaxation using the nuclear norm regularizer [26]. In [2], Askari et al. consider neural net objectives which are convex over blocks of variables. A number of recent results considered the gradient descent method on the non-convex training objective, and proved that it recovers the planted model parameters under distributional assumptions on the training data [10, 23, 25]. We refer the reader to [12, 13, 11] for other theoretical results regarding learning ReLU units and convolutional nets.

Our contributions can be summarized as follows. First, we propose a randomized convex relaxation of learning single hidden layer neural networks in the original parameter space. This should be contrasted with [26], where the convex program is in the lifted space of matrices, and is not guaranteed to find the global optimum. Our derivation also explains why direct relaxations fail and a randomized perturbation is needed. Second, we prove that the relaxation obtains the global optimum under a planted model assumption with Gaussian training data with certain probability. Our results are geometric in nature, and has a close connection to the phase transitions in compressed sensing [5, 7, 8, 9, 22, 21], escape from a mesh phenomena in Gaussian random matrices [14, 1] and sketching [18, 19, 20, 24]. Third, our numerical results highlight a phase transition, where the global optimum can be recovered when the number of samples exceeds a threshold that depends on the dimension of the filter and the number of hidden neurons. Our approach provides a general framework for obtaining convex relaxations which can be used in other architectures such as soft-max classifiers, autoencoders and recurrent nets.

2. RELAXATIONS FOR A SINGLE NEURON

Consider the problem of fitting a single neuron to predict a continuous labels y from observations x . Suppose that we observe the training data $\{x_i\}_{i=1}^n$ and labels $\{y_i\}_{i=1}^n$.

$$p^* = \min_{w \in \mathbb{R}^d} \sum_{i=1}^n (\sigma(x_i^T w) - y_i)^2, \quad (1)$$

where $\sigma(\cdot)$ is the activation function. In general, the optimization of this objective for an arbitrary training set is computationally intractable. We refer the reader to recent works on the NP-hardness of ReLU regression and approximation algorithms [6]. We rewrite the objective in (1) by introducing an additional slack variable as follows

$$p^* = \min_{w \in \mathbb{R}^d, z \in \mathbb{R}^n} \sum_{i=1}^n (z_i - y_i)^2, \quad (2)$$

$$z_i = \sigma(x_i^T w), i = 1, \dots, n \quad (3)$$

and consequently we relax the equality constraint into an inequality constraint and obtain

$$p^* \geq p = \min_{w, z} \sum_{i=1}^n (z_i - y_i)^2. \quad (4)$$

$$z_i \geq \sigma(x_i^T w), i = 1, \dots, n \quad (5)$$

The above problem is a finite dimensional convex optimization problem which can be solved very efficiently [3].

2.1. Failure of the naive relaxation

As a result of the relaxation, the convex optimization problem in (5) may not be a satisfactory approximation of the original problem, and may have more than one optimal solution. Let us illustrate the case for the ReLU activation $\sigma(u) = (u)_+$. The convex relaxation (5) becomes

$$p_{ReLU} = \min_{w, z} \sum_{i=1}^n (z_i - y_i)^2 \quad (6)$$

$$\begin{aligned} z_i &\geq (x_i^T w)_+, i = 1, \dots, n \\ &= \min_{w, z} \sum_{i=1}^n (z_i - y_i)^2, \quad (7) \\ z_i &\geq x_i^T w, i = 1, \dots, n \\ z_i &\geq 0, i = 1, \dots, n \end{aligned}$$

where we have used the fact that $(u)_+ \leq z$ holds if and only if $u \leq z$ and $0 \leq z$. Observe that $w = 0$ and $z = y$ is feasible in the constraint set of (7), and also minimizes the convex objective $\|y - z\|_2^2$. Hence, the pair $w = 0, z = y$ belongs to the set of optimal solutions of (7), regardless of the data x_1, \dots, x_n and labels y_1, \dots, y_n . Let us define w^* as the solution to (1), and suppose that the optimal value is zero. Note that the pair $z = y = (Xw^*)_+$ and $w = w^*$ also belongs to the set of optimal solutions. Therefore, the set of optimal solutions to the convex program (7) is not a singleton, and always contains the trivial solution, along with the optimal solution to (1).

It is surprising that even in the idealized case where the labels are generated by $y = (Xw^*)_+$, and features are i.i.d. Gaussian distributed, i.e., $x_i \sim \mathcal{N}(0, I_d)$, $i = 1, \dots, n$, the

above relaxation fails to recover the correct weight vector w^* . In contrast, it is possible to recover w^* using a very simple procedure as long as $n > 2d$, and n is large enough. We can consider the subset of labels which are strictly positive, $S = \bigcup_{i \in \{1, \dots, n\} : y_i > 0} \{i\}$, and attempt to solve for w^* using the pseudoinverse via $\hat{w} = X_S^\dagger y_S$. It's straightforward to show that $\hat{w} = w^*$ as long as $|S| \geq d$ and X_S has full column rank. For i.i.d. Gaussian features, this holds with high probability when n is large enough.

2.2. Randomized Convex Relaxation

In convex relaxations, relaxing the equality constraints into inequality constraints can lead to multiple spurious feasible points. It is clear that with the naive convex relaxation, we can't hope to recover the optimal solution to (1). In this section we will propose a randomly perturbed convex program aiming to recover the solution of the original problem. The reasoning behind optimizing a random objective is to randomly pick a solution among multiple feasible solutions. We will pick a random vector r distributed as $\mathcal{N}(0, I_d)$ and solve

$$\begin{aligned} \hat{w} &= \arg \min_{w \in \mathbb{R}^d, z \in \mathbb{R}^n} \frac{1}{2} \|z - y\|_2^2 + \beta r^T w, \quad (8) \\ z_i &\geq \sigma(x_i^T w), i = 1, \dots, n, \end{aligned}$$

where $\beta > 0$ is a small regularization parameter that controls the amount of random perturbation. For the ReLU activation, the above is a second-order cone program which can be solved efficiently [3]. The next theorem shows that a small random perturbation allows exact recovery of the global optimum as $\beta \rightarrow 0$.

Theorem 1. *Let $\sigma(u) = \max(u, 0)$, the ReLU activation, $x_i \sim \mathcal{N}(0, I_d)$, $i = 1, \dots, n$ and the global minimum value of (1) is $p^* = 0$, and is achieved¹ by w^* . Then, the solution of (8) as $\beta \rightarrow 0$ is equal to the global minimizer w^* with probability $\frac{1}{2} - c_2 e^{-c_3 d}$ if $n \geq c_1 d$, where c_1, c_2, c_3 are universal constants.*

This theorem essentially implies that a random perturbation to the objective is equally likely to return $w = w^*$ or $w = 0$, which are the only extreme points asymptotically as $n \rightarrow \infty$. *Proof sketch:* Plugging in the ReLU activation function $\sigma(u) = (u)_+$ we can express the optimization problem as $\beta \rightarrow 0$ via a linear program

$$\begin{aligned} \min \quad & r^T w \\ \text{s.t.} \quad & y = z, Xw \leq z, 0 \leq z. \end{aligned}$$

The pair $z = y = (Xw^*)_+$ and $w = w^*$ are optimal if

$$(y - z) + r = r \in \text{cone } X_S.$$

¹In other words, the response y is generated by a network such that $y = \sigma(x_i^T w^*)$, $i = 1, \dots, n$ holds. This is a common assumption which is also used by many others in the literature, and sometimes referred as the teacher network assumption.

The dual linear program is given by

$$\begin{aligned} \min_{u \geq 0} \quad & y^T u \\ X^T u &= r. \end{aligned}$$

Optimality conditions for the primal-dual pair are as follows

$$\exists u : u \geq 0, u_{S^c} = 0, X^T u = r,$$

where $S = \bigcup_{i \in \{1, \dots, n\} : x_i^T w^* > 0} \{i\}$ is the subset of indices over $\{1, \dots, n\}$ where the ReLU is active. This condition can be equivalently represented as a cone intersection

$$r \in \text{cone}(X_S^T),$$

where X_S is the submatrix of X composed of the rows that are in S . Without loss of generality, let us take $w^* = e_1$ due to the rotation invariance of the i.i.d. Gaussian distribution over the features. In this case, we have $S = \{i : x_{i1} > 0\}$ and S is a random set independent of x_{ij} for $j \neq 0$. Partitioning r and X_S into $r = [r_1; r']$ and $X_S^T = [x_{1S}^T; X_S'^T]$, the cone condition reduces to $r_1 = x_{1S}^T u_S$, and $r' = (X_S')^T u_S$. Fixing $r_1 > 0$ and x_{1S} , the above condition can be stated as the nullspace-cone intersection $[u_S; 1] \in \text{Null}([X_S' - r])$ and $[x_{1S}^T; -r_1]^T [u_S; 1] = 0, u_S \geq 0$. Since $[X_S' - r]$ is an i.i.d. Gaussian matrix, the nullspace-cone intersection probability can be lower-bounded by calculating the Gaussian width of the cone $([u_S; 1] : [x_{1S}^T; -r_1]^T [u_S; 1] = 0, u_S \geq 0)$ [1]. For $r_1 > 0$, a calculation of the Gaussian width yields $O(d)$, which implies that for $n \geq c_1 d$, the probability of cone intersection is $1 - c_2 e^{-c_3 d}$ conditioned on $r_1 \geq 0$. Noting that $P(r_1 > 0) = \frac{1}{2}$, we obtain the claimed result.

3. CONVOLUTIONAL NETS AND MULTIPLE NEURONS

We now consider multiple neurons where each neuron receives the output of a convolution. When the filter w is applied in a non-overlapping fashion, we need to solve the non-convex problem $\min_{w, c} \sum_{i=1}^n (y_i - \sum_{j=1}^k c_j (X_{ij} w)_+)^2$, where $X_{ij} \in \mathbb{R}^{1 \times d/k}$ is the j 'th $\frac{d}{k}$ -length block of the i 'th row of X . If we assume that $c \geq 0$, then without loss of generality, we can instead solve the problem $\min_w \sum_{i=1}^n (y_i - \sum_{j=1}^k (X_{ij} w)_+)^2$. This follows from the fact that $c_j (X_{ij} w)_+ = (c_j X_{ij} w)_+$ for $c_j \geq 0$ and therefore it is possible to implicitly include the parameter c_j in the parameter w . We relax this non-convex problem the same way we did for the single neuron case and obtain the relaxed linear program as we let $\beta \rightarrow 0$

$$\begin{aligned} \max_{z_{ij} \geq 0, w} \quad & r^T w \\ \sum_{j=1}^k z_{ij} &= y_i, i = 1, \dots, n \\ z_{ij} &\geq X_{ij} w, \forall (i, j) \in \{1, \dots, n\} \times \{1, \dots, k\}. \end{aligned} \quad (9)$$

The dual of (9) is given by

$$\begin{aligned} \min_{\lambda_{ij} \geq 0, v} \quad & y^T v \\ \lambda_{ij} &\leq v_i, \forall (i, j) \in \{1, \dots, n\} \times \{1, \dots, k\} \\ \sum_{i=1}^n \sum_{j=1}^k X_{ij}^T \lambda_{ij} &= r, \end{aligned} \quad (10)$$

where $\lambda \in \mathbb{R}^{n \times k}$ and $v \in \mathbb{R}^n$. Defining the sets S_j for $j = 1, \dots, k$ as $S_j = \bigcup_{i \in \{1, \dots, n\} : X_{ij} w^* > 0} \{i\}$, we can rewrite the equality constraint of (9) as

$$\sum_{j=1}^k \sum_{i \in S_j} X_{ij}^T \lambda_{ij} + \sum_{j=1}^k \sum_{i \in S_j^c} X_{ij}^T \lambda_{ij} = r.$$

Substituting λ_{ij} with $v_i + (\lambda_{ij} - v_i)$ and multiplying both sides by w^* , we obtain

$$\begin{aligned} \sum_{j=1}^k \sum_{i \in S_j} X_{ij} w^* v_i + \sum_{j=1}^k \sum_{i \in S_j} X_{ij} w^* (\lambda_{ij} - v_i) \\ + \sum_{j=1}^k \sum_{i \in S_j^c} X_{ij} w^* \lambda_{ij} = r^T w^*. \end{aligned} \quad (11)$$

Note that the first term of the LHS of (11) is equal to the objective of the dual (this follows from the definition that $\text{ReLU}(X_{ij} w^*) = 0$ for $i \notin S_j$), and the RHS is equal to the optimal value of the primal. For optimality, we must have the second and the third terms of the LHS to be zero since they are both nonpositive. This implies that $\lambda_{ij} = 0$ for all $i \notin S_j$, and $\lambda_{ij} = v_i$ for all $i \in S_j$.

Hence, the optimality condition is that there exists $v \in \mathbb{R}^n$ such that

$$\sum_{j=1}^k X_{S_j, j}^T v_{S_j} = r, \text{ and } v_{S_j} \geq 0, j = 1, \dots, k, \quad (12)$$

where v_{S_j} is a $|S_j|$ -dimensional vector with entries from v corresponding to the indices S_j . To reach an equivalent condition to (12) that will be useful in our analysis, let us define the sets R_i for each sample $i = 1, \dots, n$, which indicate the indices j for which the i 'th sample satisfies $X_{ij} w^* > 0$:

$$R_i = \bigcup_{j \in \{1, \dots, k\} : i \in S_j} \{j\} \quad (13)$$

The condition (12) is satisfied if and only if r falls inside the cone defined by the vectors $\sum_{j \in R_i} X_{ij}^T, i = 1, \dots, n$, that is,

$$r \in \text{cone} \left(\bigcup_{i \in \{1, \dots, n\}} \left\{ \sum_{j \in R_i} X_{ij}^T \right\} \right). \quad (14)$$

Now, we are ready to show that (14) is satisfied with probability $\frac{1}{2} - c_2 e^{-c_3 d}$ if $n \geq c_1 d$, where c_1, c_2 , and c_3 are constants (not necessarily the same constants for the single neuron case in Theorem 1). Assuming X and r both have i.i.d. $\mathcal{N}(0, 1)$ entries, and observing that we can consider only the samples with $|R_i| = 1$, we can use the same reasoning in Theorem 1. The expected number of samples with $|R_i| = 1$ is $n \frac{1}{2} (\frac{1}{2})^{k-1} = \frac{n}{2^k}$. For a fixed k , the number of samples with $|R_i| = 1$ is linear in n , and the constant term $\frac{1}{2^k}$ may as well be hidden in the constant c_1 . Now using the same argument from Theorem 1, it is straightforward to show that the success probability is $\frac{1}{2} - c_2 e^{-c_3 d}$ for $n \geq c_1 d$.

4. NUMERICAL RESULTS

In this section, we present the results of numerical experiments. The first subsection serves to visualize and compare the phase transition phenomena for our proposed relaxation method and gradient descent. In all of the experiments, the filter is applied with no overlaps, thus the filter size is equal to d/k . The second subsection presents numerical experiments on the MNIST dataset.

4.1. Phase Transition Plots for Gaussian Distribution

For all the phase transition plots in Fig. 1, for given n and d , we generate a random data matrix X , and a random filter w^* and compute the output y without adding noise (i.e., teacher network assumption). Then we run the proposed relaxation method and gradient descent algorithms separately and repeat it 100 times for each method. The plots show the minimum of the resulting 100 l_2 -norm errors between w^* and \hat{w} , that is $\|w^* - \hat{w}\|_2$, using scaled colors. Fig. 1 illustrates the performances of the proposed relaxation method and gradient descent for $k = 1, 2, 5$ neurons when the distribution of the input is Gaussian, $\mathcal{N}(0, I_d)$. Fig. 1 shows that as k increases, the probabilities of recovering w^* go up for both methods. For a given k , gradient descent seems to outperform the proposed relaxation method as the line where the phase transition occurs has a higher slope. However, gradient descent does not offer the same flexibility the proposed method does since the proposed relaxation is a convex problem and can handle extra convex constraints. Fig. 1 also shows that as k increases, the difference between the performances of the convex relaxation and gradient descent vanishes. We believe that this implies that the non-convex loss surface of the convolutional nets is becoming more like a convex surface as k increases.

4.2. Experiments on MNIST Dataset

We now present the experiment results on a randomly rotated version of the MNIST dataset [17], where the task is to predict the rotation angles of handwritten digits. The results are given in Table 1. We compare two methods where we perform least squares (LS) with l_2 regularization, on different features. For the baseline we use the original pixels, and in the second method we augment the original pixels with the filtered features where the filter is obtained by fitting the proposed re-

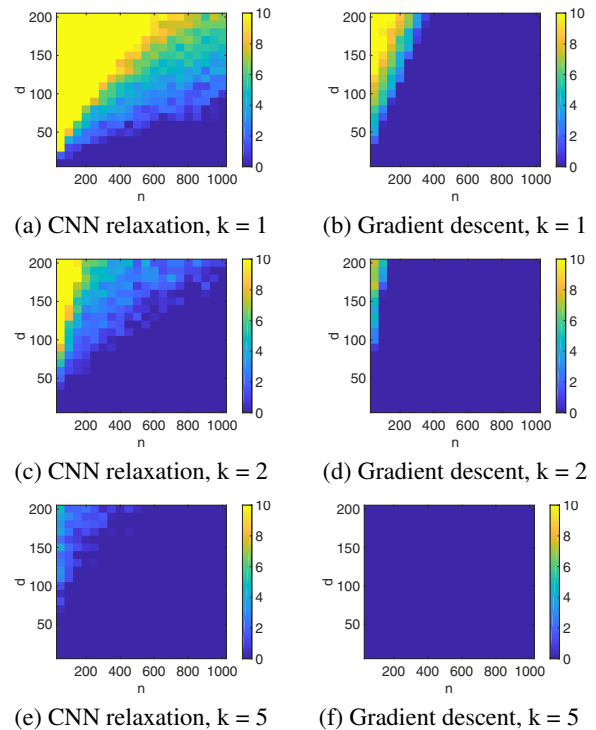


Fig. 1. Phase transition plots comparing CNN relaxation and gradient descent when the training data is i.i.d. Gaussian.

laxation to y . The results in Table 1 show root-mean-square

Table 1. MNIST experiment results.

Experiment	RMSE
LS using raw pixels	17.04
LS with learned filter	16.59

errors (RMSE) for the predicted rotation angles. These results have been obtained on the test set, which has 5000 images unseen to the training process. Table 1 shows that the proposed relaxation method helps extract useful features that make the model generalize better to the test data.

5. CONCLUSION

We investigated convex relaxations for single hidden layer no-overlap convolutional neural nets. We proved that under the planted model assumption, the relaxation method finds the global optimum with probability $\frac{1}{2}$. It is possible to make the success probability arbitrarily close to 1 by running the algorithm multiple times. We gave phase transition plots to help identify the behavior of the proposed convex relaxation in different parameter regimes. We also presented a numerical study on a real dataset and empirically showed that the proposed convex relaxation is able to extract useful features. We believe this work provides insights towards understanding the relationship between the non-convex nature of deep learning and convex relaxations.

6. REFERENCES

- [1] Dennis Amelunxen, Martin Lotz, Michael B. McCoy, and Joel A. Tropp. Living on the edge: A geometric theory of phase transitions in convex optimization. Technical report, DTIC Document, 2013.
- [2] Armin Askari, Geoffrey Negiar, Rajiv Sambharya, and Laurent El Ghaoui. Lifted neural networks. *arXiv preprint arXiv:1805.01532*, 2018.
- [3] Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004.
- [4] Alon Brutzkus and Amir Globerson. Globally optimal gradient descent for a convnet with gaussian inputs. *arXiv preprint arXiv:1702.07966*, 2017.
- [5] Emmanuel Candes and Terence Tao. Decoding by linear programming. *IEEE Trans. Info Theory*, 51(12):4203–4215, December 2005.
- [6] Santanu S. Dey, Guanyi Wang, and Yao Xie. Relu regression: Complexity, exact and approximation algorithms. *arXiv preprint arXiv:1810.03592*, 2018.
- [7] David L. Donoho. Compressed sensing. *IEEE Trans. Info. Theory*, 52(4):1289–1306, April 2006.
- [8] David L. Donoho, Iain Johnstone, and Andrea Montanari. Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising. *IEEE Trans. Info. Theory*, 59(6):3396 – 3433, 2013.
- [9] David L. Donoho and Jared Tanner. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 367(1906):4273–4293, 2009.
- [10] Simon S. Du, Jason D. Lee, Yuandong Tian, Barnabas Póczos, and Aarti Singh. Gradient descent learns one-hidden-layer cnn: Don’t be afraid of spurious local minima. *arXiv preprint arXiv:1712.00779*, 2017.
- [11] Simon S. Du, Yining Wang, Xiyu Zhai, Sivaraman Balakrishnan, Ruslan Salakhutdinov, and Aarti Singh. How many samples are needed to learn a convolutional neural network? *arXiv preprint arXiv:1805.07883*, 2018.
- [12] Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the relu in polynomial time. *arXiv preprint arXiv:1611.10258*, 2016.
- [13] Surbhi Goel, Adam Klivans, and Raghu Meka. Learning one convolutional layer with overlapping patches. *arXiv preprint arXiv:1802.02547*, 2018.
- [14] Yehoram Gordon. On milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n . In *Geometric Aspects of Functional Analysis*, pages 84–106. Springer, 1988.
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [16] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [18] Mert Pilanci and Martin J. Wainwright. Randomized sketches of convex programs with sharp guarantees. *IEEE Trans. Info. Theory*, 9(61):5096–5115, Sept 2015.
- [19] Mert Pilanci and Martin J. Wainwright. Iterative hessian sketch: Fast and accurate solution approximation for constrained least-squares. *The Journal of Machine Learning Research*, 17(1):1842–1879, 2016.
- [20] Mert Pilanci and Martin J. Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, 27(1):205–245, 2017.
- [21] Mert Pilanci, Martin J. Wainwright, and Laurent El Ghaoui. Sparse learning via boolean relaxations. *Mathematical Programming*, 151(1):63–87, 2015.
- [22] Mark Rudelson and Roman Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 61(8):1025–1045, 2008.
- [23] Mahdi Soltanolkotabi. Learning relus via gradient descent. In *Advances in Neural Information Processing Systems*, pages 2007–2017, 2017.
- [24] Yun Yang, Mert Pilanci, Martin J. Wainwright, et al. Randomized sketches for kernels: Fast and optimal nonparametric regression. *The Annals of Statistics*, 45(3):991–1023, 2017.
- [25] Xiao Zhang, Yaodong Yu, Lingxiao Wang, and Quanguan Gu. Learning one-hidden-layer relu networks via gradient descent. *arXiv preprint arXiv:1806.07808*, 2018.
- [26] Yuchen Zhang, Percy Liang, and Martin J. Wainwright. Convexified convolutional neural networks. *arXiv preprint arXiv:1609.01000*, 2016.