
Learning ReLUs via Gradient Descent

Mahdi Soltanolkotabi

Ming Hsieh Department of Electrical Engineering
University of Southern California
Los Angeles, CA
soltanol@usc.edu

Abstract

In this paper we study the problem of learning Rectified Linear Units (ReLUs) which are functions of the form $x \mapsto \max(0, \langle \mathbf{w}, \mathbf{x} \rangle)$ with $\mathbf{w} \in \mathbb{R}^d$ denoting the weight vector. We study this problem in the high-dimensional regime where the number of observations are fewer than the dimension of the weight vector. We assume that the weight vector belongs to some closed set (convex or nonconvex) which captures known side-information about its structure. We focus on the realizable model where the inputs are chosen i.i.d. from a Gaussian distribution and the labels are generated according to a planted weight vector. We show that projected gradient descent, when initialized at $\mathbf{0}$, converges at a linear rate to the planted model with a number of samples that is optimal up to numerical constants. Our results on the dynamics of convergence of these very shallow neural nets may provide some insights towards understanding the dynamics of deeper architectures.

1 Introduction

Nonlinear data-fitting problems are fundamental to many supervised learning tasks in signal processing and machine learning. Given training data consisting of n pairs of input features $\mathbf{x}_i \in \mathbb{R}^d$ and desired outputs $y_i \in \mathbb{R}$ we wish to infer a function that best explains the training data. In this paper we focus on fitting Rectified Linear Units (ReLUs) to the data which are functions $\phi_{\mathbf{w}} : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form

$$\phi_{\mathbf{w}}(\mathbf{x}) = \max(0, \langle \mathbf{w}, \mathbf{x} \rangle).$$

A natural approach to fitting ReLUs to data is via minimizing the least-squares misfit aggregated over the data. This optimization problem takes the form

$$\min_{\mathbf{w} \in \mathbb{R}^d} \mathcal{L}(\mathbf{w}) := \frac{1}{n} \sum_{i=1}^n (\max(0, \langle \mathbf{w}, \mathbf{x}_i \rangle) - y_i)^2 \quad \text{subject to} \quad \mathcal{R}(\mathbf{w}) \leq R, \quad (1.1)$$

with $\mathcal{R} : \mathbb{R}^d \rightarrow \mathbb{R}$ denoting a regularization function that encodes prior information on the weight vector.

Fitting nonlinear models such as ReLUs have a rich history in statistics and learning theory [12] with interesting new developments emerging [6] (we shall discuss all these results in greater detail in Section 5). Most recently, nonlinear data fitting problems in the form of neural networks (a.k.a. deep learning) have emerged as powerful tools for automatically extracting interpretable and actionable information from raw forms of data, leading to striking breakthroughs in a multitude of applications [13, 15, 4]. In these and many other empirical domains it is common to use local search heuristics such as gradient or stochastic gradient descent for nonlinear data fitting. These local search heuristics are surprisingly effective on real or randomly generated data. However, despite their empirical success the reasons for their effectiveness remains mysterious.

Focusing on fitting ReLUs, a-priori it is completely unclear why local search heuristics such as gradient descent should converge for problems of the form (1.1), as not only the regularization function maybe nonconvex but also the loss function! Efficient fitting of ReLUs in this high-dimensional setting poses new challenges: When are the iterates able to escape local optima and saddle points and converge to global optima? How many samples do we need? How does the number of samples depend on the a-priori prior knowledge available about the weights? What regularizer is best suited to utilizing a particular form of prior knowledge? How many passes (or iterations) of the algorithm is required to get to an accurate solution? At the heart of answering these questions is the ability to predict convergence behavior/rate of (non)convex constrained optimization algorithms. In this paper we build up on a new framework developed in the context of phase retrieval [21] for analyzing nonconvex optimization problems to address such challenges.

2 Precise measures for statistical resources

We wish to characterize the rates of convergence for the projected gradient updates (3.2) as a function of the number of samples, the available prior knowledge and the choice of the regularizer. To make these connections precise and quantitative we need a few definitions. Naturally the required number of samples for reliable data fitting depends on how well the regularization function \mathcal{R} can capture the properties of the weight vector \mathbf{w} . For example, if we know that the weight vector is approximately sparse, naturally using an ℓ_1 norm for the regularizer is superior to using an ℓ_2 regularizer. To quantify this capability we first need a couple of standard definitions which we adapt from [17, 18, 21].

Definition 2.1 (Descent set and cone) *The set of descent of a function \mathcal{R} at a point \mathbf{w}^* is defined as*

$$\mathcal{D}_{\mathcal{R}}(\mathbf{w}^*) = \left\{ \mathbf{h} : \mathcal{R}(\mathbf{w}^* + \mathbf{h}) \leq \mathcal{R}(\mathbf{w}^*) \right\}.$$

The cone of descent is defined as a closed cone $\mathcal{C}_{\mathcal{R}}(\mathbf{w}^)$ that contains the descent set, i.e. $\mathcal{D}_{\mathcal{R}}(\mathbf{w}^*) \subset \mathcal{C}_{\mathcal{R}}(\mathbf{w}^*)$. The tangent cone is the conic hull of the descent set. That is, the smallest closed cone $\mathcal{C}_{\mathcal{R}}(\mathbf{w}^*)$ obeying $\mathcal{D}_{\mathcal{R}}(\mathbf{w}^*) \subset \mathcal{C}_{\mathcal{R}}(\mathbf{w}^*)$.*

We note that the capability of the regularizer \mathcal{R} in capturing the properties of the unknown weight vector \mathbf{w}^* depends on the size of the descent cone $\mathcal{C}_{\mathcal{R}}(\mathbf{w}^*)$. The smaller this cone is the more suited the function \mathcal{R} is at capturing the properties of \mathbf{w}^* . To quantify the size of this set we shall use the notion of mean width.

Definition 2.2 (Gaussian width) *The Gaussian width of a set $\mathcal{C} \in \mathbb{R}^d$ is defined as:*

$$\omega(\mathcal{C}) := \mathbb{E}_{\mathbf{g}} \left[\sup_{\mathbf{z} \in \mathcal{C}} \langle \mathbf{g}, \mathbf{z} \rangle \right],$$

where the expectation is taken over $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. Throughout we use $\mathcal{B}^d / \mathbb{S}^{d-1}$ to denote the the unit ball/sphere of \mathbb{R}^d .

We now have all the definitions in place to quantify the capability of the function \mathcal{R} in capturing the properties of the unknown parameter \mathbf{w}^* . This naturally leads us to the definition of the minimum required number of samples.

Definition 2.3 (minimal number of samples) *Let $\mathcal{C}_{\mathcal{R}}(\mathbf{w}^*)$ be a cone of descent of \mathcal{R} at \mathbf{w}^* . We define the minimal sample function as*

$$\mathcal{M}(\mathcal{R}, \mathbf{w}^*) = \omega^2(\mathcal{C}_{\mathcal{R}}(\mathbf{w}^*) \cap \mathcal{B}^d).$$

We shall often use the short hand $n_0 = \mathcal{M}(\mathcal{R}, \mathbf{w}^)$ with the dependence on $\mathcal{R}, \mathbf{w}^*$ implied.*

We note that n_0 is exactly the minimum number of samples required for structured signal recovery from linear measurements when using convex regularizers [3, 1]. Specifically, the optimization problem

$$\sum_{i=1}^n (y_i - \langle \mathbf{x}_i, \mathbf{w} \rangle)^2 \quad \text{subject to} \quad \mathcal{R}(\mathbf{w}) \leq \mathcal{R}(\mathbf{w}^*), \quad (2.1)$$

succeeds at recovering an unknown weight vector \mathbf{w}^* with high probability from n observations of the form $\mathbf{y}_i = \langle \mathbf{a}_i, \mathbf{w}^* \rangle$ if and only if $n \geq n_0$.¹ While this result is only known to be true for convex regularization functions we believe that n_0 also characterizes the minimal number of samples even for nonconvex regularizers in (2.1). See [17] for some results in the nonconvex case as well as the role this quantity plays in the computational complexity of projected gradient schemes for linear inverse problems. Given that with nonlinear samples we have less information (we lose some information compared to linear observations) we can not hope to recover the weight vector from $n \leq n_0$ when using (1.1). Therefore, we can use n_0 as a lower-bound on the minimum number of observations required for projected gradient descent iterations (3.2) to succeed at finding the right model.

3 Theoretical results for learning ReLUs

A simple heuristic for optimizing (1.1) is to use gradient descent. One challenging aspect of the above loss function is that it is not differentiable and it is not clear how to run projected gradient descent. However, this does not pose a fundamental challenge as the loss function is differentiable except for isolated points and we can use the notion of generalized gradients to define the gradient at a non-differentiable point as one of the limit points of the gradient in a local neighborhood of the non-differentiable point. For the loss in (1.1) the generalized gradient takes the form

$$\nabla \mathcal{L}(\mathbf{w}) := \frac{1}{n} \sum_{i=1}^n (\text{ReLU}(\langle \mathbf{w}, \mathbf{x}_i \rangle) - y_i) (1 + \text{sgn}(\langle \mathbf{w}, \mathbf{x}_i \rangle)) \mathbf{x}_i. \quad (3.1)$$

Therefore, projected gradient descent takes the form

$$\mathbf{w}_{\tau+1} = \mathcal{P}_{\mathcal{K}}(\mathbf{w}_{\tau} - \mu_{\tau} \nabla \mathcal{L}(\mathbf{w}_{\tau})), \quad (3.2)$$

where μ_{τ} is the step size and $\mathcal{K} = \{\mathbf{w} \in \mathbb{R}^d : \mathcal{R}(\mathbf{w}) \leq R\}$ is the constraint set with $\mathcal{P}_{\mathcal{K}}$ denoting the Euclidean projection onto this set.

Theorem 3.1 *Let $\mathbf{w}^* \in \mathbb{R}^d$ be an arbitrary weight vector and $\mathcal{R} : \mathbb{R}^d \rightarrow \mathbb{R}$ be a proper function (convex or nonconvex). Suppose the feature vectors $\mathbf{x}_i \in \mathbb{R}^d$ are i.i.d. Gaussian random vectors distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I})$ with the corresponding labels given by*

$$\mathbf{y}_i = \max(0, \langle \mathbf{x}_i, \mathbf{w}^* \rangle).$$

To estimate \mathbf{w}^ , we start from the initial point $\mathbf{w}_0 = \mathbf{0}$ and apply the Projected Gradient Descent (PGD) updates of the form*

$$\mathbf{w}_{\tau+1} = \mathcal{P}_{\mathcal{K}}(\mathbf{w}_{\tau} - \mu_{\tau} \nabla \mathcal{L}(\mathbf{w}_{\tau})), \quad (3.3)$$

with $\mathcal{K} := \{\mathbf{w} \in \mathbb{R}^d : \mathcal{R}(\mathbf{w}) \leq \mathcal{R}(\mathbf{w}^)\}$ and $\nabla \mathcal{L}$ defined via (3.1). Also set the learning parameter sequence to $\mu_0 = 2$ and $\mu_{\tau} = 1$ for all $\tau = 1, 2, \dots$ and let $n_0 = \mathcal{M}(\mathcal{R}, \mathbf{w}^*)$, per Definition 2.3, be our lower bound on the number of observations. Also assume*

$$n > cn_0, \quad (3.4)$$

holds for a fixed numerical constant c . Then there is an event of probability at least $1 - 9e^{-\gamma n}$ such that on this event the updates (3.3) obey

$$\|\mathbf{w}_{\tau} - \mathbf{w}^*\|_{\ell_2} \leq \left(\frac{1}{2}\right)^{\tau} \|\mathbf{w}^*\|_{\ell_2}. \quad (3.5)$$

Here γ is a fixed numerical constant.

The first interesting and perhaps surprising aspect of this result is its generality: it applies not only to convex regularization functions but also nonconvex ones! As we mentioned earlier the optimization problem in (1.1) is not known to be tractable even for convex regularizers. Despite the nonconvexity of both the objective and regularizer, the theorem above shows that with a near minimal number

¹We would like to note that n_0 only approximately characterizes the minimum number of samples required. A more precise characterization is $\phi^{-1}(\omega^2(\mathcal{C}_{\mathcal{R}}(\mathbf{w}^*) \cap \mathcal{B}^d)) \approx \omega^2(\mathcal{C}_{\mathcal{R}}(\mathbf{w}^*) \cap \mathcal{B}^d)$ where $\phi(t) = \sqrt{2} \frac{\Gamma(\frac{t+1}{2})}{\Gamma(\frac{t}{2})} \approx \sqrt{t}$. However, since our results have unspecified constants we avoid this more accurate characterization.




Figure 1: Estimation error ($\|\mathbf{w}_\tau - \mathbf{w}^*\|_{\ell_2}$) obtained via running PGD iterates as a function of the number of iterations τ . The plots are for two different observations models: 1) ReLU observations of the form $\mathbf{y} = \text{ReLU}(\mathbf{X}\mathbf{w}^*)$ and 2) linear observations of the form $\mathbf{y} = \mathbf{X}\mathbf{w}^*$. The bold colors depict average behavior over 100 trials. None bold color depict the estimation error of some sample trials.

of data samples, projected gradient descent provably learns the original weight vector \mathbf{w}^* without getting trapped in any local optima.

Another interesting aspect of the above result is that the convergence rate is linear. Therefore, to achieve a relative error of ϵ the total number of iterations is on the order of $\mathcal{O}(\log(1/\epsilon))$. Thus the overall computational complexity is on the order of $\mathcal{O}(nd \log(1/\epsilon))$ (in general the cost is the total number of iterations multiplied by the cost of applying the feature matrix \mathbf{X} and its transpose). As a result, the computational complexity is also now optimal in terms of dependence on the matrix dimensions. Indeed, for a dense matrix even verifying that a good solution has been achieved requires one matrix-vector multiplication which takes $\mathcal{O}(nd)$ time.

4 Numerical experiments

In this section we carry out a simple numerical experiment to corroborate our theoretical results. For this purpose we generate a unit norm sparse vector $\mathbf{w}^* \in \mathbb{R}^d$ of dimension $d = 1000$ containing $s = d/50$ non-zero entries. We also generate a random feature matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ with $n = \lfloor 8s \log(d/s) \rfloor$ and containing i.i.d. $\mathcal{N}(0, 1)$ entries. We now take two sets of observations of size n from $\boldsymbol{\theta}^*$:

- ReLU observations: the response vector is equal to $\mathbf{y} = \text{ReLU}(\mathbf{X}\mathbf{w}^*)$.
- Linear observations: the response is $\mathbf{y} = \mathbf{X}\mathbf{w}^*$.

We apply the projected gradient iterations to both observation models starting from $\mathbf{w}_0 = \mathbf{0}$. For the ReLU observations we use the step size discussed in Theorem 3.1. For the linear model we apply projected gradient descent updates of the form

$$\mathbf{w}_{\tau+1} = \mathcal{P}_{\mathcal{K}} \left(\mathbf{w}_\tau - \frac{1}{n} \mathbf{X}^T (\mathbf{X}\mathbf{w}_\tau - \mathbf{y}) \right).$$

In both cases we use the regularizer $\mathcal{R}(\mathbf{w}) = \|\mathbf{w}\|_{\ell_0}$ so that the projection only keeps the top s entries of the vector (a.k.a. iterative hard thresholding). In Figure 1 the resulting estimation errors ($\|\mathbf{w}_\tau - \mathbf{w}^*\|_{\ell_2}$) is depicted as a function of the number of iterations τ . The bold colors depict average behavior over 100 trials. The estimation error of some sample trials are also depicted in none bold

colors. This plot clearly show that PGD iterates applied to ReLU observations converge quickly to the ground truth. This figure also clearly demonstrates that the behavior of the PGD iterates applied to both models are similar, further corroborating the results of Theorem 3.1. We note that the sample complexity used in this simulation is $8s \log(n/s)$ which is a constant factor away from $n_0 \propto s \log(n/s)$ confirming our assertion that the required sample complexity is a constant factor away from n_0 (as predicted by Theorem 3.1).

5 Discussions and prior art

There is a large body of work on learning nonlinear models. A particular class of such problems that have been studied are the so called idealized Single Index Models (SIMs) [9, 10]. In these problems the inputs are labeled examples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \in \mathbb{R}^d \times \mathbb{R}$ which are guaranteed to satisfy $y_i = f(\langle \mathbf{w}, \mathbf{x}_i \rangle)$ for some $\mathbf{w} \in \mathbb{R}^d$ and nondecreasing (Lipchitz continuous) $f : \mathbb{R} \rightarrow \mathbb{R}$. The goal in this problem is to find a (nearly) accurate such f and \mathbf{w} . An interesting polynomial-time algorithm called the Isotron exists for this problem [12, 11]. In principle, this approach can also be used to fit ReLUs. However, these results differ from ours in term of both assumptions and results. On the one had, the assumptions are slightly more restrictive as they require bounded features \mathbf{x}_i , outputs y_i and weights. On the other hand, these result hold for much more general distributions and more general models than the realizable model studied in this paper. These results also do not apply in the high dimensional regime where the number of observations is significantly smaller than the number of parameters (see [5] for some results in this direction). In the realizable case, the Isotron result require $\mathcal{O}(\frac{1}{\epsilon})$ iterations to achieve ϵ error in objective value. In comparison, our results guarantee convergence to a solution with relative error ϵ ($\|\mathbf{w}_\tau - \mathbf{w}^*\|_{\ell_2} / \|\mathbf{w}^*\|_{\ell_2} \leq \epsilon$) after $\log(1/\epsilon)$ iterations. Focusing on the specific case of ReLU functions, an interesting recent result [6] shows that reliable learning of ReLUs is possible under very general but bounded distributional assumptions. To achieve an accuracy of ϵ the algorithm runs in $\text{poly}(1/\epsilon)$ time. In comparison, as mentioned earlier our result requires $\log(1/\epsilon)$ iterations for reliable parameter estimation. We note however we study the problem in different settings and a direct comparison is not possible between the two results.

We would like to note that there is an interesting growing literature on learning shallow neural networks with a single hidden layer with i.i.d. inputs, and under a realizable model (i.e. the labels are generated from a network with planted weights) [23, 2, 25]. For isotropic Gaussian inputs, [23] shows that with two hidden unites ($k = 2$) there are no critical points for configurations where both weight vectors fall into (or outside) the cone of ground truth weights. With the same assumptions, [2] proves that for a single-hidden ReLU network with a single non-overlapping convolutional filter, all local minimizers of the population loss are global; they also give counter-examples in the overlapping case and prove the problem is NP-hard when inputs are not Gaussian. [25] studies general single-hidden layer networks and shows that a version of gradient descent which uses a fresh batch of samples in each iteration converges to the planted model. This holds using an initialization obtained via a tensor decomposition method. Our approach and convergence results differ from this literature in a variety of different ways. First, we focus on zero hidden layers with a regularization term. Some of this literature focuses on one-hidden layers without (or with specific) regularization. Second, unlike some of these results such as [2, 14], we study the optimization properties of the empirical function, not its expected value. Third, we initialize at zero in lieu of sophisticated initialization schemes. Finally, our framework does not require a fresh batch of samples per new gradient iteration as in [25]. We also note that several publications study the effect of over-parametrization on the training of neural networks without any regularization [19, 8, 16, 22]. Therefore, the global optima are not unique and hence the solutions may not generalize. In comparison we study the problem with an arbitrary regularization which allows for a unique global optima.

6 Proofs

6.1 Preliminaries

In this section we gather some useful results on concentration of stochastic processes which will be crucial in our proofs. These results are mostly adapted from [21]. We begin with a lemma which is a direct consequence of Gordon’s escape from the mesh lemma [7].

Lemma 6.1 Assume $\mathcal{C} \subset \mathbb{R}^d$ is a cone and \mathbb{S}^{d-1} is the unit sphere of \mathbb{R}^d . Also assume that

$$n \geq \max \left(20 \frac{\omega^2(\mathcal{C} \cap \mathbb{S}^{d-1})}{\delta^2}, \frac{1}{2\delta} - 1 \right),$$

for a fixed numerical constant c . Then for all $\mathbf{h} \in \mathcal{C}$

$$\left| \frac{1}{n} \sum_{i=1}^n (\langle \mathbf{x}_i, \mathbf{h} \rangle)^2 - \|\mathbf{h}\|_{\ell_2}^2 \right| \leq \delta \|\mathbf{h}\|_{\ell_2}^2,$$

holds with probability at least $1 - 2e^{-\frac{\delta^2}{360}n}$.

We also need a generalization of the above lemma stated below.

Lemma 6.2 ([21]) Assume $\mathcal{C} \subset \mathbb{R}^d$ is a cone (not necessarily convex) and \mathbb{S}^{d-1} is the unit sphere of \mathbb{R}^d . Also assume that

$$n \geq \max \left(80 \frac{\omega^2(\mathcal{C} \cap \mathbb{S}^{d-1})}{\delta^2}, \frac{2}{\delta} - 1 \right),$$

for a fixed numerical constant c . Then for all $\mathbf{u}, \mathbf{h} \in \mathcal{C}$

$$\left| \frac{1}{n} \sum_{i=1}^n \langle \mathbf{x}_i, \mathbf{u} \rangle \langle \mathbf{x}_i, \mathbf{h} \rangle - \mathbf{u}^* \mathbf{h} \right| \leq \delta \|\mathbf{u}\|_{\ell_2} \|\mathbf{h}\|_{\ell_2},$$

holds with probability at least $1 - 6e^{-\frac{\delta^2}{1440}n}$.

We next state a generalization of Gordon's escape through the mesh lemma also from [21].

Lemma 6.3 ([21]) Let $\mathbf{s} \in \mathbb{R}^d$ be fixed vector with nonzero entries and construct the diagonal matrix $\mathbf{S} = \text{diag}(\mathbf{s})$. Also, let $\mathbf{X} \in \mathbb{R}^{n \times d}$ have i.i.d. $\mathcal{N}(0, 1)$ entries. Furthermore, assume $\mathcal{T} \subset \mathbb{R}^d$ and define $b_d(\mathbf{s}) = \mathbb{E}[\|\mathbf{S}\mathbf{g}\|_{\ell_2}]$, where $\mathbf{g} \in \mathbb{R}^d$ is distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. Also, define

$$\sigma(\mathcal{T}) := \max_{\mathbf{v} \in \mathcal{T}} \|\mathbf{v}\|_{\ell_2}.$$

Then for all $\mathbf{u} \in \mathcal{T}$

$$\left| \|\mathbf{S}\mathbf{A}\mathbf{u}\|_{\ell_2} - b_d(\mathbf{s}) \|\mathbf{u}\|_{\ell_2} \right| \leq \|\mathbf{s}\|_{\ell_\infty} \omega(\mathcal{T}) + \eta,$$

holds with probability at least $1 - 6e^{-\frac{\eta^2}{8\|\mathbf{s}\|_{\ell_\infty}^2 \sigma^2(\mathcal{T})}}$.

The previous lemma leads to the following Corollary.

Corollary 6.4 Let $\mathbf{s} \in \mathbb{R}^d$ be fixed vector with nonzero entries and assume $\mathcal{T} \subset \mathbb{B}^d$. Furthermore, assume

$$\|\mathbf{s}\|_{\ell_2}^2 \geq \max \left(20 \|\mathbf{s}\|_{\ell_\infty}^2 \frac{\omega^2(\mathcal{T})}{\delta^2}, \frac{3}{2\delta} - 1 \right).$$

Then for all $\mathbf{u} \in \mathcal{T}$,

$$\left| \frac{\sum_{i=1}^n s_i^2 (\langle \mathbf{x}_i, \mathbf{u} \rangle)^2}{\|\mathbf{s}\|_{\ell_2}^2} - \|\mathbf{u}\|_{\ell_2}^2 \right| \leq \delta,$$

holds with probability at least $1 - 6e^{-\frac{\delta^2}{1440} \|\mathbf{s}\|_{\ell_2}^2}$.

6.2 Convergence proof (Proof of Theorem 3.1)

In this section we shall prove Theorem 3.1. Throughout, we use the shorthand \mathcal{C} to denote the descent cone of \mathcal{R} at \mathbf{w}^* , i.e. $\mathcal{C} = \mathcal{C}_{\mathcal{R}}(\mathbf{w}^*)$. We begin by analyzing the first iteration. Using $\mathbf{w}_0 = \mathbf{0}$ we have

$$\mathbf{w}_1 := \mathcal{P}_{\mathcal{K}}(\mathbf{w}_0 - \mu_0 \nabla \mathcal{L}(\mathbf{w}_0)) = \mathcal{P}_{\mathcal{K}} \left(\frac{2}{n} \sum_{i=1}^n y_i \mathbf{x}_i \right) = \mathcal{P}_{\mathcal{K}} \left(\frac{2}{n} \sum_{i=1}^n \text{ReLU}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \mathbf{x}_i \right).$$

We use the argument of [21][Page 25, inequality (7.34)] which shows that

$$\|\mathbf{w}_1 - \mathbf{w}^*\|_{\ell_2} \leq 2 \cdot \sup_{\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^d} \mathbf{u}^T \left(\frac{2}{n} \sum_{i=1}^n \text{ReLU}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \mathbf{x}_i - \mathbf{w}^* \right). \quad (6.1)$$

Using $\text{ReLU}(z) = \frac{z+|z|}{2}$ we have

$$\frac{2}{n} \sum_{i=1}^n \text{ReLU}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{u} \rangle - \langle \mathbf{u}, \mathbf{w}^* \rangle = \mathbf{u}^T \left(\frac{1}{n} \mathbf{X}^T \mathbf{X} - \mathbf{I} \right) \mathbf{w}^* + \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle. \quad (6.2)$$

We proceed by bounding the first term in the above equality. To this aim we decompose \mathbf{u} in the direction parallel/perpendicular to that of \mathbf{w}^* and arrive at

$$\begin{aligned} \mathbf{u}^T \left(\frac{1}{n} \mathbf{X}^T \mathbf{X} - \mathbf{I} \right) \mathbf{w}^* &= \frac{(\mathbf{u}^T \mathbf{w}^*)}{\|\mathbf{w}^*\|_{\ell_2}^2} (\mathbf{w}^*)^T \left(\frac{1}{n} \mathbf{X}^T \mathbf{X} - \mathbf{I} \right) \mathbf{w}^* + \frac{1}{n} \left\langle \mathbf{X} \left(\mathbf{I} - \frac{\mathbf{w}^* (\mathbf{w}^*)^T}{\|\mathbf{w}^*\|_{\ell_2}^2} \right) \mathbf{u}, \mathbf{X} \mathbf{w}^* \right\rangle, \\ &\sim (\mathbf{u}^T \mathbf{w}^*) \left(\frac{\|\mathbf{g}\|_{\ell_2}^2}{n} - 1 \right) + \frac{\|\mathbf{w}^*\|_{\ell_2}}{\sqrt{n}} \mathbf{a}^T \left(\mathbf{I} - \frac{\mathbf{w}^* (\mathbf{w}^*)^T}{\|\mathbf{w}^*\|_{\ell_2}^2} \right) \mathbf{u}, \\ &\leq \|\mathbf{w}^*\|_{\ell_2} \left| \frac{\|\mathbf{g}\|_{\ell_2}^2}{n} - 1 \right| + \frac{\|\mathbf{w}^*\|_{\ell_2}}{\sqrt{n}} \sup_{\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^d} \mathbf{a}^T \left(\mathbf{I} - \frac{\mathbf{w}^* (\mathbf{w}^*)^T}{\|\mathbf{w}^*\|_{\ell_2}^2} \right) \mathbf{u}, \end{aligned} \quad (6.3)$$

with $\mathbf{g} \in \mathbb{R}^n$ and $\mathbf{a} \in \mathbb{R}^d$ are independent random Gaussian random vectors distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. By concentration of Chi-squared random variables

$$\left| \|\mathbf{g}\|_{\ell_2}^2 / n - 1 \right| \leq \Delta, \quad (6.4)$$

holds with probability at least $1 - 2e^{-n \frac{\Delta^2}{8}}$. Also,

$$\frac{1}{\sqrt{n}} \mathbf{a}^T \left(\mathbf{I} - \frac{\mathbf{w}^* (\mathbf{w}^*)^T}{\|\mathbf{w}^*\|_{\ell_2}^2} \right) \mathbf{u} \leq \frac{1}{\sqrt{n}} (\omega(\mathcal{C} \cap \mathcal{B}^d) + \eta), \quad (6.5)$$

holds with probability at least $1 - e^{-\frac{\eta^2}{2}}$. Plugging (6.4) with $\Delta = \frac{\delta}{6}$ and (6.5) with $\eta = \frac{\delta}{6} \sqrt{n}$ into (6.3), as long as $n \geq \frac{36}{\delta^2} \omega^2(\mathcal{C} \cap \mathcal{B}^d)$, then

$$\sup_{\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^d} \mathbf{u}^T \left(\frac{1}{n} \mathbf{X}^T \mathbf{X} - \mathbf{I} \right) \mathbf{w}^* \leq \frac{\delta}{2} \|\mathbf{w}^*\|_{\ell_2}, \quad (6.6)$$

holds with probability at least $1 - 3e^{-n \frac{\delta^2}{288}}$.

We now focus on bounding the second term in (6.2). To this aim we decompose \mathbf{u} in the direction parallel/perpendicular to that of \mathbf{w}^* and arrive at

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle \right| &= \left| (\mathbf{u}^T \mathbf{w}^*) \frac{1}{n} \sum_{i=1}^n \frac{|\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} + \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u}_\perp \rangle \right|, \\ &\leq \|\mathbf{w}^*\|_{\ell_2} \left| \frac{1}{n} \sum_{i=1}^n \frac{|\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} \right| + \left| \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u}_\perp \rangle \right|. \end{aligned} \quad (6.7)$$

with $\mathbf{u}_\perp = \left(\mathbf{I} - \frac{\mathbf{w}^* (\mathbf{w}^*)^T}{\|\mathbf{w}^*\|_{\ell_2}^2} \right) \mathbf{u}$. Now note that $\frac{|\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2}$ is sub-exponential and

$$\left\| \frac{|\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} \right\|_{\psi_1} \leq c,$$

with fixed numerical constant. Thus by Bernstein's type inequality ([24][Proposition 5.16])

$$\left| \frac{1}{n} \sum_{i=1}^n \frac{|\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} \right| \leq t, \quad (6.8)$$

holds with probability at least $1 - 2e^{-\gamma n \min(t^2, t)}$ with γ a fixed numerical constant.. Also note that

$$\frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u}_\perp \rangle \sim \sqrt{\frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle|^2} \frac{1}{\sqrt{n}} \langle \mathbf{g}, \mathbf{u}_\perp \rangle.$$

Furthermore, $\frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle|^2 \leq (1 + \Delta) \|\mathbf{w}^*\|_{\ell_2}^2$, holds with probability at least $1 - 2e^{-n \frac{\Delta^2}{8}}$ and

$$\sup_{\mathbf{u} \in \mathcal{C} \cap \mathbb{S}^{d-1}} |\langle \mathbf{g}, \mathbf{u}_\perp \rangle| \leq (2\omega(\mathcal{C} \cap \mathbb{S}^{d-1}) + \eta),$$

holds with probability at least $1 - e^{-\frac{\eta^2}{2}}$. Combining the last two inequalities we conclude that

$$\left| \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u}_\perp \rangle \right| \leq \sqrt{1 + \Delta} \frac{(2\omega(\mathcal{C} \cap \mathbb{S}^{d-1}) + \eta)}{\sqrt{n}} \|\mathbf{w}^*\|_{\ell_2}, \quad (6.9)$$

holds with probability at least $1 - 2e^{-n \frac{\Delta^2}{8}} - e^{-\frac{\eta^2}{2}}$. Plugging (6.8) and (6.9) with $t = \frac{\delta}{6}$, $\Delta = 1$, and $\eta = \frac{\delta}{6\sqrt{2}}\sqrt{n}$ into (6.7)

$$\left| \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle \right| \leq \frac{\delta}{2} \|\mathbf{w}^*\|_{\ell_2}, \quad (6.10)$$

holds with probability at least $1 - 3e^{-\gamma n \delta^2} - 2e^{-\frac{n}{8}}$ as long as $n \geq 288 \frac{\omega^2(\mathcal{C} \cap \mathbb{S}^{d-1})}{\delta^2}$. Thus pluggin (6.6) and (6.10) into (6.1) we conclude that for $\delta = 7/400$

$$\|\mathbf{w}_1 - \mathbf{w}^*\|_{\ell_2} \leq 2 \cdot \sup_{\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^d} \mathbf{u}^T \left(\frac{2}{n} \sum_{i=1}^n \text{ReLU}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \mathbf{x}_i - \mathbf{w}^* \right) \leq 2\delta \|\mathbf{w}^*\|_{\ell_2} \leq \frac{7}{200} \|\mathbf{w}^*\|_{\ell_2},$$

holds with probability at least $1 - 8e^{-\gamma n}$ as long as $n \geq c\omega^2(\mathcal{C} \cap \mathbb{S}^{d-1})$ for a fixed numerical constant c . To introduce our general convergence analysis we begin by defining

$$E(\epsilon) = \{ \mathbf{w} \in \mathbb{R}^d : \mathcal{R}(\mathbf{w}) \leq \mathcal{R}(\mathbf{w}^*), \|\mathbf{w} - \mathbf{w}^*\|_{\ell_2} \leq \epsilon \|\mathbf{w}^*\|_{\ell_2} \} \quad \text{with } \epsilon = \frac{7}{200}.$$

To prove Theorem 3.1 we use [21][Page 25, inequality (7.34)] which shows that if we apply the projected gradient descent update $\mathbf{w}_{\tau+1} = \mathcal{P}_{\mathcal{K}}(\mathbf{w}_\tau - \nabla \mathcal{L}(\mathbf{w}_\tau))$, the error $\mathbf{h}_\tau = \mathbf{w}_\tau - \mathbf{w}^*$ obeys

$$\|\mathbf{h}_{\tau+1}\|_{\ell_2} = \|\mathbf{w}_{\tau+1} - \mathbf{w}^*\|_{\ell_2} \leq 2 \cdot \sup_{\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^n} \mathbf{u}^* (\mathbf{h}_\tau - \nabla \mathcal{L}(\mathbf{w}_\tau)). \quad (6.11)$$

To complete the convergence analysis it is then sufficient to prove

$$\sup_{\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^n} \mathbf{u}^* (\mathbf{h}_\tau - \nabla \mathcal{L}(\mathbf{w}_\tau)) \leq \frac{1}{4} \|\mathbf{h}_\tau\|_{\ell_2} = \frac{1}{4} \|\mathbf{w}_\tau - \mathbf{w}^*\|_{\ell_2}. \quad (6.12)$$

We will instead prove that the following stronger result holds for all $\mathbf{u} \in \mathcal{C} \cap \mathcal{B}^n$ and $\mathbf{w} \in E(\epsilon)$

$$\mathbf{u}^* (\mathbf{w} - \mathbf{w}^* - \nabla \mathcal{L}(\mathbf{w})) \leq \frac{1}{4} \|\mathbf{w} - \mathbf{w}^*\|_{\ell_2}. \quad (6.13)$$

The equation (6.13) above implies (6.12) which when combined with (6.11) proves the convergence result of the Theorem (specifically equation (3.5)). The rest of this section is dedicated to proving (6.13). To this aim note that $\text{ReLU}(\langle \mathbf{x}_i, \mathbf{w} \rangle) = \frac{\langle \mathbf{x}_i, \mathbf{w} \rangle + |\langle \mathbf{x}_i, \mathbf{w} \rangle|}{2}$. Thus (see the extended version of this paper [20] for more detailed derivation of the identity below)

$$\begin{aligned} \langle \nabla \mathcal{L}(\mathbf{w}), \mathbf{u} \rangle &= \frac{1}{n} \sum_{i=1}^n \langle \mathbf{x}_i, \mathbf{w} - \mathbf{w}^* \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle + \frac{1}{n} \sum_{i=1}^n \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{w} - \mathbf{w}^* \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle \\ &\quad + \frac{1}{n} \sum_{i=1}^n (\text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) \langle \mathbf{x}_i, \mathbf{w} - \mathbf{w}^* \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle \\ &\quad + \frac{1}{2n} \sum_{i=1}^n (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) (\text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle)) |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle \end{aligned}$$

Now defining $\mathbf{h} = \mathbf{w} - \mathbf{w}^*$ we conclude that $\langle \mathbf{u}, \mathbf{w} - \mathbf{w}^* - \nabla \mathcal{L}(\mathbf{w}) \rangle = \langle \mathbf{u}, \mathbf{h} - \nabla \mathcal{L}(\mathbf{w}) \rangle$ is equal to

$$\begin{aligned} \langle \mathbf{u}, \mathbf{h} - \nabla \mathcal{L}(\mathbf{w}) \rangle &= \mathbf{u}^T \left(\mathbf{I} - \frac{1}{n} \mathbf{X} \mathbf{X}^T \right) \mathbf{h} - \frac{1}{n} \sum_{i=1}^n \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{h} \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle, \\ &+ \frac{\langle \mathbf{h}, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} \frac{1}{n} \sum_{i=1}^n (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{h} \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle, \\ &+ \frac{\text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle)}{2n} \sum_{i=1}^n (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) \\ &\quad |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle. \end{aligned}$$

Now define $\mathbf{h}_\perp = \mathbf{h} - (\mathbf{h}^T \mathbf{w}^*) / (\|\mathbf{w}^*\|_{\ell_2}^2) \mathbf{w}^*$. Using this we can rewrite the previous expression in the form (see the proof in the extended version of this paper [20] for more detailed derivation)

$$\begin{aligned} \langle \mathbf{u}, \mathbf{w} - \mathbf{w}^* - \nabla \mathcal{L}(\mathbf{w}) \rangle &= \mathbf{u}^T \left(\mathbf{I} - \frac{1}{n} \mathbf{X} \mathbf{X}^T \right) \mathbf{h} - \frac{1}{n} \sum_{i=1}^n \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{h} \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle, \\ &+ \frac{1}{n} \sum_{i=1}^n (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{h}_\perp \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle, \\ &+ \frac{1}{n} \sum_{i=1}^n \left[\frac{\text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle)}{2} (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) + \frac{\langle \mathbf{h}, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} \right] \\ &\quad (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle \end{aligned} \quad (6.14)$$

We now proceed by stating bounds on each of the four terms in (6.14). The detailed derivation of these bounds appear in the the extended version of this paper [20].

Lemma 6.5 *Assume the setup of Theorem 3.1. Then as long as $n \geq cn_0$, we have*

$$\mathbf{u}^* \left(\mathbf{I} - \frac{1}{n} \mathbf{X}^* \mathbf{X} \right) \mathbf{h} \leq \delta \|\mathbf{h}\|_{\ell_2}, \quad (6.15)$$

$$-\frac{1}{n} \sum_{i=1}^n \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{h} \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle \leq \delta \|\mathbf{h}\|_{\ell_2}, \quad (6.16)$$

$$\frac{1}{n} \sum_{i=1}^n (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) \langle \mathbf{x}_i, \mathbf{h}_\perp \rangle \langle \mathbf{x}_i, \mathbf{u} \rangle \leq 2\sqrt{1+\delta} \left(\delta + \sqrt{\frac{21}{20}} \epsilon \right) \|\mathbf{h}\|_{\ell_2}, \quad (6.17)$$

$$\begin{aligned} &\frac{1}{n} \sum_{i=1}^n \left[\frac{\text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle)}{2} (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) + \frac{\langle \mathbf{h}, \mathbf{w}^* \rangle}{\|\mathbf{w}^*\|_{\ell_2}^2} \right] \\ &\quad (1 - \text{sgn}(\langle \mathbf{x}_i, \mathbf{w} \rangle) \text{sgn}(\langle \mathbf{x}_i, \mathbf{w}^* \rangle)) |\langle \mathbf{x}_i, \mathbf{w}^* \rangle| \langle \mathbf{x}_i, \mathbf{u} \rangle \leq \frac{4\sqrt{1+\delta}}{(1-\epsilon)^2} \left(\delta + \sqrt{\frac{21}{20}} \epsilon \right) \|\mathbf{h}\|_{\ell_2}, \end{aligned} \quad (6.18)$$

holds for all $\mathbf{u} \in \mathcal{C} \cap \mathbb{S}^{d-1}$ and $\mathbf{w} \in E(\epsilon)$ with probability at least $1 - 9e^{-\gamma n}$.

Combining (6.15), (6.16), (6.17), and (6.18) we conclude that

$$\langle \mathbf{u}, \mathbf{w} - \mathbf{w}^* - \nabla \mathcal{L}(\mathbf{w}) \rangle \leq 2 \left(\delta + \sqrt{1+\delta} \left(1 + \frac{2}{(1-\epsilon)^2} \right) \left(\delta + \sqrt{\frac{21}{20}} \epsilon \right) \right) \|\mathbf{w} - \mathbf{w}^*\|_{\ell_2},$$

holds for all $\mathbf{u} \in \mathcal{C} \cap \mathbb{S}^{d-1}$ and $\mathbf{w} \in E(\epsilon)$ with probability at least $1 - 16e^{-\gamma \delta^2 n} - (n+10)e^{-\gamma n}$. Using this inequality with $\delta = 10^{-4}$ and $\epsilon = 7/200$ we conclude that $\langle \mathbf{u}, \mathbf{w} - \mathbf{w}^* - \nabla \mathcal{L}(\mathbf{w}) \rangle \leq \frac{1}{4} \|\mathbf{w} - \mathbf{w}^*\|_{\ell_2}$, holds for all $\mathbf{u} \in \mathcal{C} \cap \mathbb{S}^{d-1}$ and $\mathbf{w} \in E(\epsilon)$ with high probability.

Acknowledgements

This work was done in part while the author was visiting the Simon's Institute for the Theory of Computing. M.S. would like to thank Adam Klivans and Matus Telgarsky for discussions related to [6] and the Isotron algorithm.

References

- [1] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference*, 2014.
- [2] A. Brutzkus and A. Globerson. Globally optimal gradient descent for a convnet with gaussian inputs. *International Conference on Machine Learning (ICML)*, 2017.
- [3] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.
- [4] R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167. ACM, 2008.
- [5] R. Ganti, N. Rao, R. M. Willett, and R. Nowak. Learning single index models in high dimensions. *arXiv preprint arXiv:1506.08910*, 2015.
- [6] S. Goel, V. Kanade, A. Klivans, and J. Thaler. Reliably learning the ReLU in polynomial time. *arXiv preprint arXiv:1611.10258*, 2016.
- [7] Y. Gordon. *On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n* . Springer, 1988.
- [8] B. D. Haeffele and R. Vidal. Global optimality in tensor factorization, deep learning, and beyond. *arXiv preprint arXiv:1506.07540*, 2015.
- [9] J. L. Horowitz and W. Hardle. Direct semiparametric estimation of single-index models with discrete covariates. *Journal of the American Statistical Association*, 91(436):1632–1640, 1996.
- [10] H. Ichimura. Semiparametric least squares (SLS) and weighted SLS estimation of single-index models. *Journal of Econometrics*, 58(1-2):71–120, 1993.
- [11] S. M. Kakade, V. Kanade, O. Shamir, and A. Kalai. Efficient learning of generalized linear and single index models with isotonic regression. In *Advances in Neural Information Processing Systems*, pages 927–935, 2011.
- [12] A. T. Kalai and R. Sastry. The isotron algorithm: High-dimensional isotonic regression. In *COLT*, 2009.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [14] Y. Li and Y. Yuan. Convergence analysis of two-layer neural networks with ReLU activation. *arXiv preprint arXiv:1705.09886*, 2017.
- [15] A. Mohamed, G. E. Dahl, and G. Hinton. Acoustic modeling using deep belief networks. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(1):14–22, 2012.
- [16] Quynh Nguyen and Matthias Hein. The loss surface of deep and wide neural networks. *arXiv preprint arXiv:1704.08045*, 2017.
- [17] S. Oymak, B. Recht, and M. Soltanolkotabi. Sharp time–data tradeoffs for linear inverse problems. *arXiv preprint arXiv:1507.04793*, 2015.
- [18] S. Oymak and M. Soltanolkotabi. Fast and reliable parameter estimation from nonlinear observations. *arXiv preprint arXiv:1610.07108*, 2016.
- [19] T. Poston, C-N. Lee, Y. Choie, and Y. Kwon. Local minima and back propagation. In *Neural Networks, 1991., IJCNN-91-Seattle International Joint Conference on*, volume 2, pages 173–176. IEEE, 1991.
- [20] M. Soltanolkotabi. Learning ReLUs via gradient descent. *arXiv preprint arXiv:1705.04591*, 2017.

- [21] M. Soltanolkotabi. Structured signal recovery from quadratic measurements: Breaking sample complexity barriers via nonconvex optimization. *arXiv preprint arXiv:1702.06175*, 2017.
- [22] M. Soltanolkotabi, A. Javanmard, and J. D. Lee. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. 07 2017.
- [23] Y. Tian. An analytical formula of population gradient for two-layered relu network and its applications in convergence and critical point analysis. *International Conference on Machine Learning (ICML)*, 2017.
- [24] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [25] K. Zhong, Z. Song, P. Jain, P. L. Bartlett, and I. S. Dhillon. Recovery guarantees for one-hidden-layer neural networks. *arXiv preprint arXiv:1706.03175*, 2017.