

# Understanding How Over-Parametrization Leads to Acceleration: A case of learning a single teacher neuron

**Jun-Kun Wang**  
*Yale University*

JUN-KUN.WANG@YALE.EDU

**Jacob Abernethy**  
*Georgia Institute of Technology*

PROF@GATECH.EDU

**Editors:** Vineeth N Balasubramanian and Ivor Tsang

## Abstract

Over-parametrization has become a popular technique in deep learning. It is observed that by over-parametrization, a larger neural network needs a fewer training iterations than a smaller one to achieve a certain level of performance — namely, over-parametrization leads to acceleration in optimization. However, despite that over-parametrization is widely used nowadays, little theory is available to explain the acceleration due to over-parametrization. In this paper, we propose understanding it by studying a simple problem first. Specifically, we consider the setting that there is a single teacher neuron with quadratic activation, where over-parametrization is realized by having multiple student neurons learn the data generated from the teacher neuron. We provably show that over-parametrization helps the iterate generated by gradient descent to enter the neighborhood of a global optimal solution that achieves zero testing error faster.

**Keywords:** Over-parametrization

## 1. Introduction

Over-parametrization has become a popular technique in deep learning, as it is now widely observed larger neural nets can achieve better performance. Furthermore, a larger network can be trained to achieve a certain level of prediction performance with fewer iterations than that of a smaller net. This observation, to our knowledge, can be dated back as early as the work of [Livni et al. \(2014\)](#), who try different levels of over-parametrization and report that SGD converges much faster and finds a better solution when it is used to train a larger network. However, the reason why over-parametrization can lead to an acceleration still remains a mystery, and very little theory has helped explain the observation, with perhaps the notable exception of ([Arora et al., 2018](#)). [Arora et al. \(2018\)](#) consider over-parametrizing a single-output linear regression with  $l_p$  loss for  $p > 2$ —the square loss corresponds to  $p = 2$ —and they study the linear regression problem by replacing the model  $w \in \mathcal{R}^d$  by another model  $w_1 \in \mathcal{R}^d$  times a scalar  $w_2 \in \mathcal{R}$ . They show that the dynamics of gradient descent on the new over-parametrized model are equivalent to the dynamics of gradient descent on the original objective function with an adaptive learning rate plus some momentum terms. However, in practice, people actually use the techniques of over-parametrization, adaptive learning rate, and momentum simultaneously in deep learning (see e.g. ([Hoffer et al., 2017](#); [Kingma and Ba, 2015](#))), as each technique appears to contribute to performance and they




Figure 1: In the experiment, we set the dimension  $d = 10$  and the number of training samples  $n = 200$ . Additional 200 samples are sampled and served as testing data. We let  $w_* = e_1$  with  $e_1$  being the unit vector. Each neuron  $w^{(k)} \in \mathcal{R}^d$  ( $k \in [K]$ ) of the student network is initialized by sampling from an isotropic distribution and is close to the origin (i.e.  $w_0^{(k)} \sim 0.01 \cdot N(0, I_d/d)$ ). We apply gradient descent with the same step size  $\eta = 0.001$  to train different sizes of neural networks. Each curve represents the progress of gradient descent for different  $K$ . Subfigure (a): Objective value (2) vs. iteration  $t$  on training data (testing data, respectively). For subfigure (b) and (c), please see Section 3 for the precise definition and details.

may, to some extent, be complementary. It has been suggested that over-parameterizing a model leads implicitly to an adaptive learning rate or momentum, but this does not appear to fully explain the performance improvement.

To understand the benefits of overparameterization, let us begin by studying a simple canonical problem: a single teacher neuron  $w_* \in \mathcal{R}^d$  with quadratic activation function. Specifically, the label  $y_i$  and the design vector  $x_i \sim \mathcal{N}(0, \mathcal{I}_d)$  of sample  $i$  satisfies  $y_i = (x_i^\top w_*)^2$ . Therefore, the *standard* objective function for learning the teacher neuron  $w_*$  is

$$\min_{w \in \mathcal{R}^d} f(w) := \frac{1}{4n} \sum_{i=1}^n ((x_i^\top w)^2 - y_i)^2, \quad (1)$$

where  $n$  denotes the number of samples. Problem (1) is called *phase retrieval* in signal processing literature, which has real applications in physical science such as astronomy and microscopy (see e.g. (Candés et al., 2013), (Fannjiang and Strohmer, 2020), (Shechtman et al., 2015)). There are also some specialized algorithms designed for achieving a better sample complexity or computational complexity to recover  $w_*$  modulo the unrecoverable sign. We choose this problem as a starting point of understanding acceleration due to over-parametrization. Specifically, we consider the following way to over-parametrize the original objective (1),

$$\min_{W \in \mathcal{R}^{d \times K}} f(W) := \frac{1}{4n} \sum_{i=1}^n ((x_i^\top w^{(1)})^2 + (x_i^\top w^{(2)})^2 + \dots + (x_i^\top w^{(K)})^2 - y_i)^2, \quad (2)$$

where  $w^{(j)}$  denote the  $j$ th column of the weight matrix  $W \in \mathcal{R}^{d \times K}$ . Optimizing the objective function (2) can be viewed as training a student network with  $K$  student neurons. While a  $d$  dimensional model exists which perfectly predicts the labels generated by the teacher neuron (i.e.  $\pm w_*$ ), one can consider training a much larger model instead (i.e.  $d \times K$  number of parameters). Note that if  $K = 1$ , objective (2) reduces to the original objective (1). On the other hand, a larger  $k > 1$  means a higher degree of over-parametrization.

Let us now establish, empirically, the clear advantage of over-parametrization for accelerating the learning and optimization process. We tried  $K = \{1, 3, 10\}$  number of student neurons in (2), which represent different degrees of over-parametrization, and we applied gradient descent to train the networks. Figure 1 shows the results. The empirical findings displayed on the figure are quite stark. First, we see that over-parametrization not only helps to decrease the training error faster but also decrease the testing error faster (c.f. subfigure (a) and (b)). Furthermore, the generalization error is very small since both training error and testing error approach zero. Second, regardless of the size  $K$ , a common pattern is that the dynamics of gradient descent can be divided into two stages. In the first stage, gradient descent makes little progress on decreasing the function value; while in the second stage, the iterate generated by gradient descent exhibits a linear convergence to a global solution. Specifically, by over-parametrization, gradient descent spends fewer iterations in the first stage and enters the linear convergence regime that makes the fast progress more quickly. We provide more results in Section 4.4. Specifically, we also tried different values of the step size  $\eta$  and we observed similar patterns as Figure 1 shows. Even when gradient descent uses the best step size for each model with different neurons  $K$ , we still observe that gradient descent enters the linear convergence regime faster for a larger model. Thus, the acceleration due to over-parametrization *cannot* be simply explained by that gradient descent uses a larger *effective* step size, as the effect due to parameters  $\eta$  and  $K$  is complementary in the experiment.

## 2. Related works

**Over-parametrization:** Though our work focuses on understanding why over-parametrization leads to acceleration in optimization (i.e. improving the convergence time), we also want to acknowledge some related works of understanding over-parametrization in different aspects (e.g. (Arora et al., 2019; Brutzkus and Globerson, 2019; Emschwiller et al., 2020; Goldt et al., 2019; Tian, 2020)). There is also a trend of works studying how over-parametrization changes the optimization landscape of empirical risk minimization for neural nets with quadratic activation. (e.g. (Du and Lee, 2018; Gamarnik et al., 2019; Ge et al., 2019; Kazemipour et al., 2019; Soltanolkotabi et al., 2018; Nguyen and Hein, 2017; Venturi et al., 2019; Mannelia et al., 2020)). The goals of these works are different from ours.

**Quadratic activation and matrix sensing:** The optimization landscape of problem (1) (i.e. phase retrieval) and its variants has been studied by (Davis et al., 2018; Soltanolkotabi, 2014; Sun et al., 2016; White et al., 2016), which shows that as long as the number of samples is sufficiently large, it has no spurious local optima and all the local optima are globally optimal. Chen et al. (2019) provably show that applying gradient descent with an isotropic random initialization for solving (1) leads to an optimal solution that recovers the teacher neuron  $w_*$  modulo the unrecoverable sign. In this work we show that an over-parametrized student network trained by gradient descent takes even fewer iterations to recover  $w_*$ . We also note that the optimization problem (2) can be rewritten as the form of matrix sensing (see e.g. Gunasekar et al. (2017); Li et al. (2019, 2018); Gidel et al. (2019)).

**Learning a single neuron:** Studying learning a single neuron in non-convex optimization is not new (e.g. Goel et al. (2019); Yehudai and Shamir (2020); Kakade et al. (2011); Goel et al. (2017); Kalan et al. (2019); Soltanolkotabi (2017); Mei et al. (2018); Frei et al. (2020)).

However, those works are not for showing faster convergence by over-parametrization. The goals are different.

### 3. Preliminaries

**Notations and assumptions:** We use the notation  $W^{\#K} := [w_{\#K}^{(1)}, \dots, w_{\#K}^{(K)}] \in \mathcal{R}^{d \times K}$  to represent the weights of a student network with  $K$  number of neurons. Each column  $k$  of the matrix  $W^{\#K}$  is the weight vector  $w_{\#K}^{(k)} \in \mathcal{R}^d$  that corresponds to the neuron  $k$  of the student network. Thus,  $W^{\#1} := [w_{\#1}^{(1)}] \in \mathcal{R}^d$  is the network consists of a single neuron; while for  $K > 1$ , the notation represents the weights of an over-parametrized network. In this paper, without loss of generality, we assume that  $w_* = \|w_*\|e_1 \in \mathcal{R}^d$  with  $e_1$  being the unit vector whose first element is 1. We define the parallel component  $w_{\#K,t}^{(k),\parallel}$  and the perpendicular component  $w_{\#K,t}^{(k),\perp}$  in iteration  $t$  as follows,

$$\begin{aligned} w_{\#K,t}^{(k),\parallel} &:= \langle w_{\#K,t}^{(k)}, w_* \rangle = \|w_*\| w_{\#K,t}^{(k)}[1] \\ w_{\#K,t}^{(k),\perp} &:= [w_{\#K,t}^{(k)}[2], \dots, w_{\#K,t}^{(k)}[d]]^\top. \end{aligned} \quad (3)$$

Namely, the parallel component  $w_{\#K,t}^{(k),\parallel}$  is the projection of a student neuron  $k$  learned in iteration  $t$  on the teacher neuron  $w_*$ ; while the perpendicular component  $w_{\#K,t}^{(k),\perp}$  is the  $d-1$  dimensional sub-vector of  $w_{\#K,t}^{(k)}$  excluding the first element. We assume that each neuron  $k$  of each network,  $w_{\#K,0}^{(k)} \in \mathcal{R}^d$ , is initialized i.i.d. randomly from an isotropic distribution (e.g. gaussian distribution).

**Metric of the progress in optimization:** The first challenge to show that over-parametrization leads to acceleration in optimization is the design of the metric for the comparison. Since different  $K$ 's corresponds to different optimization problems, the notion of *acceleration* here is non-standard in optimization literature. In the optimization literature, acceleration usually means that an algorithm takes a fewer iterations than other algorithms for the *same* optimization problem. Fortunately, by exploiting the problem structure, we can have a natural metric of the progress as follows. For any size of the student network  $W \in \mathcal{R}^{d \times K}$ , we consider

$$\text{dist}(W, w_*) := \min_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \|W - w_* q^\top\|. \quad (4)$$

This is due to the observation that for any  $K$ , the global optimal solutions of (2) that achieve zero testing error are  $w_* q^\top \in \mathcal{R}^{d \times K}$  for any  $q \in \mathcal{R}^K$  such that  $\|q\|_2 = 1$ . To see this, substitute  $W = w_* q^\top \in \mathcal{R}^{d \times K}$  into (2). We have that for any  $x_i \in \mathcal{R}^d$  it holds that  $(x_i^\top w^{(1)})^2 + (x_i^\top w^{(2)})^2 + \dots + (x_i^\top w^{(K)})^2 - y_i = \|x_i^\top W\|_F^2 - (x_i^\top w_*)^2 = \text{tr}((x_i^\top w_* q^\top)^\top (x_i^\top w_* q^\top)) - (x_i^\top w_*)^2 = 0$ . Therefore, the metric  $\text{dist}(W, w_*)$  as be viewed as a surrogate of the testing error. In particular,  $\text{dist}(W_t, w_*)$  represents the distance of the current iterate  $W_t$  and its closest global optimal solution to the over-parametrized objective (2) that achieves zero testing error. Note that the argmin of (4) is  $q_* := \frac{W^\top w_*}{\|W^\top w_*\|_2} = \arg \min_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \|W - w_* q^\top\|$ . On sub-figure (c) of Figure 1, we plot the distance of the iterates generated by gradient descent and its closet global optimal solution for different sizes  $K$  of neural nets. We see that over-parametrization enables shrinking the distance  $\text{dist}(W_t^{\#K}, w_*)$  faster.

**Gradient descent dynamics:** For the notation brevity, we will suppress the symbol  $\#K$  when it is clear in the context. The gradient of a student neuron  $w^{(k)}$  for the over-parametrized problem (2) is

$$\nabla_{w^{(k)}} f(W) := \frac{1}{n} \sum_{i=1}^n ((x_i^\top w^{(1)})^2 + (x_i^\top w^{(2)})^2 + \dots + (x_i^\top w^{(K)})^2 - y_i)(x_i^\top w^{(k)})x_i. \quad (5)$$

Its expectation, which is the population gradient of a student neuron  $w^{(k)}$  (i.e. gradient when the number of samples  $n$  is infinite), is

$$\begin{aligned} \mathbb{E}_{x \sim N(0, I_d)} [\nabla_{w^{(k)}} f(W)] &= (3\|w^{(k)}\|^2 - \|w_*\|^2)w^{(k)} - 2(w_*^\top w^{(k)})w_* \\ &\quad + \sum_{j \neq k}^K 2((w^{(j)})^\top w^{(k)})w^{(j)} + \|w^{(j)}\|^2 w^{(k)}, \end{aligned} \quad (6)$$

where we use the fact that for any vector  $u, v \in \mathcal{R}^d$ ,  $\mathbb{E}_{x \sim N(0, I_d)} [(x^\top u)^3 x] = 3\|u\|^2 u$  and  $\mathbb{E}_{x \sim N(0, I_d)} [(x^\top u)^2 (x^\top v)x] = 2(u^\top v)u + \|u\|^2 v$ . For  $K = 1$ , (6) becomes  $\mathbb{E}[\nabla_{w^{(1)}} f(W^{\#1})] = (3\|w^{(1)}\|^2 - \|w_*\|^2)w^{(1)} - 2(w_*^\top w^{(1)})w_*$ . If gradient descent uses the population gradient for the update (i.e.  $w_{\#1, t+1}^{(1)} = w_{\#1, t}^{(1)} - \eta \mathbb{E}[\nabla_{w^{(1)}} f(W_t^{\#1})]$ ), then the dynamics of the student network consists of a single neuron (i.e.  $K = 1$ ) evolves as follows,

$$\begin{aligned} w_{t+1}^{(1), \parallel} &= w_t^{(1), \parallel} (1 + \eta(3\|w_*\|^2 - 3\|w_t^{(1)}\|^2)) \\ w_{t+1}^{(1), \perp} &= w_t^{(1), \perp} (1 + \eta(\|w_*\|^2 - 3\|w_t^{(1)}\|^2)). \end{aligned} \quad (7)$$

On the other hand, if a student network has  $K > 1$  neurons, the dynamics of each neuron  $k$  of the student network evolves as follows,

$$\begin{aligned} w_{t+1}^{(k), \parallel} &= \underbrace{w_t^{(k), \parallel} (1 + 3\eta(\|w_*\|^2 - \|w_t^{(k)}\|^2))}_{\text{component A}} - \underbrace{(2\eta \sum_{j \neq k}^K \langle w_t^{(j)}, w_t^{(k)} \rangle w_t^{(j), \parallel} + \eta w_t^{(k), \parallel} \sum_{j \neq k}^K \|w_t^{(j)}\|^2)}_{\text{component B}} \\ w_{t+1}^{(k), \perp} &= \underbrace{w_t^{(k), \perp} (1 + \eta(\|w_*\|^2 - 3\|w_t^{(k)}\|^2))}_{\text{component C}} - \underbrace{(2\eta \sum_{j \neq k}^K \langle w_t^{(j)}, w_t^{(k)} \rangle w_t^{(j), \perp} + \eta w_t^{(k), \perp} \sum_{j \neq k}^K \|w_t^{(j)}\|^2)}_{\text{component D}} \end{aligned} \quad (8)$$

where both the component  $B$  of  $w_{t+1}^{(k), \parallel}$  and the component  $D$  of  $w_{t+1}^{(k), \perp}$  can be viewed as the terms due to the interaction of student neuron  $k$  and the other student neurons.

**More observations:** On Subfigure (c) of Figure 1, we plot a quantity over iterations, which is

$$\text{vec}(w_* q_t^\top - W_t^{\#K})^\top \nabla^2 f(W_t^{\#K}) \text{vec}(w_* q_t^\top - W_t^{\#K}), \quad (9)$$

where  $\nabla^2 f(W_t^{\#K}) \in \mathcal{R}^{dK \times dK}$  is the Hessian and  $w_* q_t^\top$  is the closet global optimal solution to  $W_t^{\#K}$  and the notation  $\text{vec}(\cdot)$  represents the vectorization operation of its matrix argument. The quantity can be viewed as a measure of the strong convexity. Specifically, if the quantity is larger than 0, then it suggests that the current optimization landscape is strongly convex with respect to  $w_* q_t^\top$ . Hence, the observation suggests that gradient descent enters a benign region faster for a larger student network. In the following section, We will answer why over-parametrization helps gradient descent to enter a region that has the benign optimization landscape faster.

## 4. Analysis

In this section, we answer why over-parametrization leads to the acceleration. We first show that gradient descent (GD) exhibits linear convergence to a global optimal solution when  $\text{dist}(W^{\#K}, w_*)$  is small. We then answer why over-parametrization helps the iterate to enter the neighborhood of a global optimal solution faster. For the ease of analysis, we assume that gradient descent uses population gradient (6) for the update and we denote  $\nabla F(W) := \mathbb{E}_x[\nabla f(W)]$  and  $\nabla^2 F(W) := \mathbb{E}_x[\nabla^2 f(W)]$  accordingly.

### 4.1. When does the iterate generated by gradient descent enter a benign region?

We first introduce a key lemma. The lemma shows that whenever the iterate is in the neighborhood of a global optimal solution, gradient descent has a linear convergence rate.

**Lemma 1** (*locally linear convergence*) *Suppose that at time  $t_0$ ,  $\text{dist}(W_{t_0}, w_*) := \|W_{t_0} - w_* q_{t_0}^\top\| \leq \nu \|w_*\|$  where  $W_{t_0} \in \mathcal{R}^{d \times K}$ ,  $\nu > 0$  satisfies  $2 - 14\nu - 2\nu^2 > 0$ , and  $q_{t_0} := \arg \min_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \|W_{t_0} - w_* q^\top\|$ . Then, gradient descent with the step size  $\eta \leq \frac{2-14\nu-2\nu^2}{(13+16\nu^2)^2 \|w_*\|^4}$  generates iterates  $\{W_t\}_{t \geq t_0}$  satisfying  $\text{dist}^2(W_{t+1}, w_*) \leq (1 - \eta(2 - 14\nu - 2\nu^2)) \text{dist}^2(W_t, w_*)$ .*

The proof is available in Appendix A. Note that the lemma holds for *any* size  $K$  of neural nets, as long as the condition,  $\text{dist}(W, w_*) \leq \nu \|w_*\|_2$  with the required  $\nu$ , is satisfied. The condition ensures the locally linear convergence of gradient descent and might be easily satisfied by having a larger number of student neurons  $K > 1$ . Specifically, each neuron  $w^{(k)} \in \mathcal{R}^d$  of  $W \in \mathcal{R}^{d \times K}$  only needs to have a smaller component on the direction of  $w_*$  in order to have  $W$  be sufficiently close to the teacher neuron  $w_*$  up to a transform  $q^\top$ , compared to the case when one only has a single student neuron ( $K = 1$ ). To support this argument, we plot the quantities  $|w_{\#K,t}^{(k)}| = |\langle w_{\#K,t}^{(k)}, w_* \rangle|$  for each  $k \in [K]$  of different student networks trained by gradient descent on Figure 2. The figure shows that with more student neurons  $K$ , each  $w^{(k)}$  only needs a smaller projection on  $w_*$  for the aggregate projection  $\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k)}|^2}$  to achieve certain level. In the later subsections, we will provide a formal analysis.

### 4.2. How does gradient descent work for the student network consists of a single neuron?

In this subsection, we analyze the case that the student network only has a single neuron. For brevity, we suppress the notation  $\#1$  in the following. Define  $T_\gamma := \min\{t : \|w_t[1]\| - \|w_*\| \leq \gamma \text{ and } \|w_t^\perp\| \leq \gamma\}$ . Note that if  $\|w_t[1]\| - \|w_*\| \leq \gamma$  and that  $\|w_t^\perp\| \leq \gamma$ , then  $\text{dist}^2(w_t, w_*) = \|w_t[1]\| - \|w_*\|^2 + \|w_t^\perp\|_2^2 \leq 2\gamma^2$ . Let us assume that (C1)  $\|w_*\| > 1.1\gamma$  (strong signal) and (C2)  $\gamma \geq 10\|w_0\|$  (small initialization). Note that to invoke Lemma 1 for showing locally linear convergence after the iterate gets into a benign region, we will set  $2\gamma^2 = \nu^2 \|w_*\|^2$  with  $\nu$  satisfying  $2 - 14\nu - 2\nu^2 > 0$  (i.e.  $\nu \leq 0.141$ ); consequently, (C1) is trivially satisfied.

**Theorem 2** *Suppose that the conditions (C1-C2) hold. Assume that the step size satisfies  $\eta \leq c/\|w_*\|^2$  for some sufficiently small constant  $c > 0$ . Then gradient descent for problem*




Figure 2: Subfigure (a) and (b): parallel component  $|w_{\#K,t}^{(k),||}|$  of each neuron  $k$  vs. iteration  $t$  for student networks with size  $K = \{3, 10\}$  trained by gradient descent. The curve “l2 norm” represents  $\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k),||}|^2}$ , which can be viewed as the aggregate projection of a student network  $W \in \mathcal{R}^{d \times K}$  on the teacher neuron  $w_* \in \mathcal{R}^d$ . For comparison, we also plot the curve for  $|w_{\#1,t}^{(1),||}|$  labeled by  $K = 1$  on the same subfigures. Subfigure (c) and (d): the square norm  $\|w_{\#K,t}^{(k)}\|^2$  of each neuron  $k$  for different  $K$ ’s. The curve “sum” represents  $\|W^{\#K}\|_F^2$ . For comparison, we also plot the curve for  $\|w_{\#1,t}^{(1)}\|^2$  labeled by  $K = 1$  on the subfigures.

(1) (i.e.  $K = 1$ ) has  $T_\gamma \leq \frac{\log(\frac{\|w_*\| - \gamma}{|w_0[1]|})}{\log(1 + \eta\Delta)}$ , where  $\Delta := 6\gamma(\|w_*\| - \gamma) > 0$ . Furthermore, for  $0 \leq t \leq T_\gamma$ , we have that  $|w_t^\parallel| \geq (1 + \eta\Delta)^t |w_0^\parallel|$  and  $\|w_t^\perp\| \leq \gamma$ .

Theorem 2 states that to achieve  $\text{dist}^2(w_t, w_*) \leq 2\gamma^2$ , gradient descent only needs at most  $\log(\frac{\|w_*\| - \gamma}{|w_0[1]|}) / \log(1 + \eta\Delta)$  number of iterations. Furthermore, on the signal direction,  $|w_t[1]|$  (and hence  $w_t^\parallel$ ) grows exponentially at a rate at least  $1 + \eta\Delta$  before reaching at  $\|w_*\| - \gamma$ . On the other hand, by a close-to-zero initialization (C2), the perpendicular component  $\|w_t^\perp\|$  remains small. Consequently, we have that  $\text{dist}^2(W_t^{\#1}, w_*) \leq \max(\gamma^2, (\|w_*\| - |w_0[1]|)(1 + \eta\Delta)^t)^2 + \gamma^2$ , for  $0 \leq t \leq T_\gamma$  before gradient descent enters the linear convergence regime. The proof of Theorem 2 is available in Appendix B. A similar result as Theorem 2 was shown in (Chen et al., 2019).



### 4.3. How does over-parametrization help entering a benign region?

Let us begin by providing a more detailed observation regarding the dynamics of gradient descent. Figure 3 plots each component of  $w_{\#K,t}^{(k),\parallel}$  and  $w_{\#K,t}^{(k),\perp}$  in the gradient descent dynamics (8), while Figure 4 plots (the magnitude of) component B to component A and the ratio of component D to component C for  $K = 10$ . The empirical findings show that the component due to the interaction (component B, or component D respectively) is negligible compared to the component that is without the dependency on the other neurons (component A, or component C respectively). Based on the observation, we can re-write the population dynamics (8) in the early stage (i.e. before gradient descent enters the linear convergence regime) as follows,

$$\begin{aligned} |w_{\#K,t+1}^{(k),\parallel}| &\geq (1 - \theta) |w_{\#K,t}^{(k),\parallel}| (1 + \eta(3\|w_*\|^2 - 3\|w_{\#K,t}^{(k),\parallel}\|^2)) \\ \|w_{\#K,t+1}^{(k),\perp}\| &\leq (1 + \vartheta) \|w_{\#K,t}^{(k),\perp}\| (1 + \eta(\|w_*\|^2 - 3\|w_{\#K,t}^{(k),\parallel}\|^2)). \end{aligned} \quad (10)$$

for some small numbers  $\theta, \vartheta \ll 1$  (empirically  $\approx 10^{-4}$  as Figure 3 shows).

That is, in the early stage, the approximated dynamics (10) of each student neuron of an over-parametrized network does not deviate too much from the original dynamics without over-parametrization (7). The approximated dynamics (10) will be only used for analyzing the stage before the iterate enters the benign region, i.e. used only before entering a linear convergence regime, though empirically the approximated dynamics hold all the time during the execution of the algorithm (Figure 3).

On the other hand, by comparing subfigure (a) and subfigure (b) of Figure 3, we see that component A and B of each neuron  $k$  are with the same sign during the execution. This observation together with the dynamics (8) tend to implies that  $|w_{\#K,t+1}^{(k),\parallel}| \leq |w_{\#K,t}^{(k),\parallel}| (1 + \eta(3\|w_*\|^2 - 3\|w_{\#K,t}^{(k),\parallel}\|^2))$  when  $K > 1$ . On the other hand, the dynamics (7) has  $|w_{\#1,t+1}^{(1),\parallel}| = |w_{\#1,t}^{(1),\parallel}| (1 + \eta(3\|w_*\|^2 - 3\|w_{\#1,t}^{(1),\parallel}\|^2))$ . Also, as the case of single neuron, the perpendicular component of each neuron  $k$  of  $K$  (i.e.  $\|w_{\#K,t}^{(k),\perp}\|$ ) remains small (Figures 5). Consequently, we could write

$$|w_{\#1,t}^{(1),\parallel}| \gtrsim |w_{\#K,t}^{(k),\parallel}| \text{ and } \|w_{\#1,t}^{(1),\perp}\| \gtrsim \|w_{\#K,t}^{(k),\perp}\|, \quad (11)$$

where the approximation  $\gtrsim$  accounts for the fact that the size of initial points due to the random initialization may be different and that the small interaction components are present in the dynamics for  $K > 1$ . Figure 2 confirms that the relation (11) generally holds on average empirically.

**Lemma 3** *Suppose that the approximated dynamics (10) and (11) hold from iteration 0 to iteration  $t$ . Then, the network with a single neuron and an over-parametrized network trained by GD with the same step size  $\eta$  has  $\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k),\parallel}|^2} \gtrsim \sqrt{(1 - 2t\theta)} \sqrt{K} |w_{\#1,t}^{(1),\parallel}|$ .*

Lemma 3 states that if the single neuron of the non-overparametrized network has a certain projection on  $w_*$  at time  $t$ , then the over-parametrized network with  $K$  neurons will have approximately  $\sqrt{K}$  times larger projection on  $w_*$ , modulo the  $\sqrt{1 - 2t\theta}$  factor which is close to 1 if the product  $t\theta$  is small (as Figure 3 shows). This demonstrates the advantage of over-parametrization — over-parametrization helps to make more progress on growing the model's projection on  $w_*$ .






Figure 3: Subfigure (a) shows  $w_t^{(k),\parallel} (1 + \eta(3\|w_*\|^2 - 3\|w_t^{(k)}\|^2))$  (i.e. component A of  $w_{\#K,t}^{(k),\parallel}$ ) versus iteration  $t$  for each neuron  $k$ . Subfigure (b) plots  $2\eta \sum_{j \neq k}^K ((w_t^{(j)})^\top w_t^{(k)}) w_t^{(j),\parallel} + \eta w_t^{(k),\parallel} \sum_{j \neq k}^K \|w_t^{(j)}\|^2$  (i.e. component B of  $w_{\#K,t}^{(k),\parallel}$ ) versus iteration  $t$  for each neuron  $k$ . Subfigure (c) plots the norm of component C of  $w_{\#K,t}^{(k),\perp}$ , while subfigure (d) plots the norm of component D of  $w_{\#K,t}^{(k),\perp}$  for each neuron  $k$ . The empirical findings show that the components due to interaction of the other neurons (i.e. component B and D) are small (notice that the scale of the vertical axis of (a) and (b), (c) and (d) are different) compared to their counterparts (i.e. component A and C respectively), which suggests that  $\theta, \vartheta \approx 10^{-4}$  on (10) empirically. Similar patterns exhibit in training under different  $K$ 's.

**Lemma 4** Suppose that  $\eta \leq \frac{1}{3\|w_*\|^2}$ . By following the conditions as Lemma 3, we have that

$$\|w_{\#K,t}^{(k),\perp}\| \lesssim \frac{|w_{\#K,t}^{(k),\parallel}| |w_{\#K,0}^{(k),\perp}|}{|w_{\#K,0}^{(k),\parallel}|} \frac{1}{\psi^t} \lesssim \frac{|w_{\#1,t}^{(1),\parallel}| |w_{\#K,0}^{(k),\perp}|}{|w_{\#K,0}^{(k),\parallel}|} \frac{1}{\psi^t}, \text{ where } \psi := (1 - \theta - \vartheta - \theta\vartheta)(1 + \eta\|w_*\|^2).$$

Lemma 4 states that the ratio of the perpendicular component  $\|w_{\#K,t}^{(k),\perp}\|$  to the parallel component  $|w_{\#K,t}^{(k),\parallel}|$  of each neuron decays exponentially if  $\psi > 1$  (which holds if  $\eta\|w_*\|^2 \gtrsim \frac{\theta + \vartheta}{1 - \theta - \vartheta}$ ). By combining Lemma 3 and 4, we have the following theorem, which characterizes the difference of the distances to  $w_*$  at iteration  $t$ .

**Theorem 5** (Snapshot at  $t$ ) Suppose that the approximated dynamics (10) and (11) hold from 0 to  $t$  and that at iteration  $t$ , the student network with a single neuron trained by GD with the step size  $\eta$  has  $|w_{\#1,t}^{(1),\parallel}| = c_{1,t}\|w_*\|^2$  for some number  $c_{1,t}$  satisfying  $1 > c_{1,t} > 0$ .

Denote  $c_{2,t}$  a number that satisfies  $\frac{c_{1,t}\|w_{\#K,0}^{(k),\perp}\|}{|w_{\#K,0}^{(k),\parallel}|} \frac{1}{\psi^t} \leq c_{2,t}$  for each  $k \in [K]$ . Suppose that the




Figure 4: The ratio of (the magnitude of) component B to component A and the ratio of component D to component C. The plot shows that the components due to the interaction (component B and component D) are negligible.




Figure 5: The perpendicular component  $\|w_{\#K,t}^{(k),\perp}\|$  of each  $k$  over iterations  $t$ . We see that the perpendicular component remains small, compared to the signal component.

step size  $\eta$  also satisfies  $\eta \leq \frac{1}{3\|w_*\|^2}$  and makes  $\psi := (1 - \theta - \vartheta - \theta\vartheta)(1 + \eta\|w_*\|^2) > 1$ . Then, an over-parametrized network  $W_t^{\#K}$  trained by GD with the same  $\eta$  has

$$\text{dist}^2(W_t^{\#1}, w_*) - \text{dist}^2(W_t^{\#K}, w_*) \gtrsim \|w_*\|^2 (2c_{1,t}(\sqrt{(1-2t\theta)}\sqrt{K} - 1) - K(c_{1,t}^2 + c_{2,t}^2) + c_{1,t}^2).$$

Recall that in Theorem 2, we upper-bound  $\text{dist}^2(W_t^{\#1}, w_*)$ . Simply combining Theorem 2 and 5 leads to a distance upper-bound of  $\text{dist}^2(W_t^{\#K}, w_*)$  at certain iteration  $t$ . The lower bound of the difference of the distances in Theorem 5 shows a strict improvement due to over-parametrization when it is positive, which answers why over-parametrization helps gradient descent to enter the linear convergence regime faster — gradient descent for an over-parametrized network shrinks the distance to  $w_*$  faster in the early stage. Note that the lower bound is a quadratic function of  $\sqrt{K}$  and is increasing for  $1 \leq \sqrt{K} \leq \frac{c_{1,t}\sqrt{1-2t\theta}}{c_{1,t}^2 + c_{2,t}^2}$ , which means that up to a certain threshold of  $K$ , more over-parametrization could lead to more improvements. Moreover, if  $c_{1,t}$  and  $c_{2,t}$  further satisfy (\*)  $2c_{1,t}(\sqrt{2}\sqrt{(1-2t\theta)} - 1) - c_{1,t}^2 - 2c_{2,t}^2 > 0$ , then the lower bound of the difference for  $K = 2$  neurons is strictly positive and keeps being positive up to  $\sqrt{K} \leq \lfloor \frac{c_{1,t}\sqrt{1-2t\theta} + \sqrt{c_{1,t}^2(1-2t\theta) - (c_{1,t}^2 + c_{2,t}^2)(2c_{1,t} - c_{1,t}^2)}}{c_{1,t}^2 + c_{2,t}^2} \rfloor$ , which

gives an upper limit of the degree of over-parametrization that allows acceleration. The condition (\*) is easily satisfied when (1)  $t\theta \ll 1$  so that  $\sqrt{(1-2t\theta)} \cong 1$  and (2)  $c_{2,t} \ll 1$ , which happens when the approximated dynamics (8) holds for a small  $\theta$  and that the ratio of the perpendicular component to the parallel component of neuron  $k$  decays sufficiently fast (Lemma 4).

**Proof** [of Theorem 5]

Let us compute  $\text{dist}(W_t^{\#1}, w_*)$  and  $\text{dist}(W_t^{\#k}, w_*)$ . For  $\text{dist}(W_t^{\#1}, w_*)^2$ , we have that

$$\begin{aligned} \text{dist}(W_t^{\#1}, w_*)^2 &= \min_{q \in \{-1, +1\}} \|W_t^{\#1} - w_* q^\top\|^2 = \|w_{\#1,t}^{(1)}\|^2 - 2|w_{\#1,t}^{(1)}| + \|w_*\|^2 \\ &= \|w_{\#1,t}^{(1)}\|^2 + \|w_*\|^2 - 2\sqrt{\|w_{\#1,t}^{(1)}\|^2}. \end{aligned} \quad (12)$$

On the other hand, for  $\text{dist}^2(W_t^{\#K}, w_*)$ , we have that

$$\begin{aligned} \text{dist}^2(W_t^{\#K}, w_*) &= \min_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \|W_t^{\#K} - w_* q^\top\|_F^2 \\ &= \min_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \text{tr}((W_t^{\#K} - w_* q^\top)^\top (W_t^{\#K} - w_* q^\top)) \\ &= \|W_t^{\#K}\|_F^2 + \|w_*\|^2 - 2 \max_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \text{tr}((W_t^{\#K})^\top w_* q^\top) \\ &= \|W_t^{\#K}\|_F^2 + \|w_*\|^2 - 2\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k)}|^2}, \end{aligned} \quad (13)$$

where the last inequality is due to that

$$\max_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \text{tr}((W_t^{\#K})^\top w_* q^\top) \stackrel{(a)}{=} \max_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \text{tr}(\bar{v} q^\top) = \max_{q \in \mathcal{R}^K: \|q\|_2 \leq 1} \text{tr}(q^\top \bar{v}) \stackrel{(b)}{=} \|\bar{v}\|. \quad (14)$$

where (a) we denote  $\bar{v} := [w_{\#K,t}^{(1)}, w_{\#K,t}^{(2)}, \dots, w_{\#K,t}^{(K)}]^\top \in \mathcal{R}^K$  and (b) is because  $q = \bar{v}/\|\bar{v}\|$ . Combining (12) and (13), we have that

$$\begin{aligned} &\text{dist}^2(W_t^{\#1}, w_*) - \text{dist}^2(W_t^{\#K}, w_*) \\ &= 2(\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k)}|^2} - \sqrt{\|w_{\#1,t}^{(1)}\|^2}) - (\|W_t^{\#K}\|_F^2 - \|W_t^{\#1}\|_F^2). \end{aligned} \quad (15)$$

To continue, we need the lower bound of  $\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k)}|^2} - \sqrt{\|w_{\#1,t}^{(1)}\|^2}$  and the upper bound of  $\|W_t^{\#K}\|_F^2 - \|W_t^{\#1}\|_F^2$ . By Lemma 3, we have that

$$\sqrt{\sum_{k=1}^K |w_{\#K,t}^{(k)}|^2} - \sqrt{\|w_{\#1,t}^{(1)}\|^2} \gtrsim (\sqrt{(1-2t\theta)}\sqrt{K}-1)\|w_{\#1,t}^{(1)}\| = (\sqrt{(1-2t\theta)}\sqrt{K}-1)c_{1,t}\|w_*\|^2. \quad (16)$$

On the other hand, for the difference of the norms, we have that  $\|W_t^{\#1}\|^2 \geq |w_{\#1,t}^{(1)}|^2 = \frac{|w_{\#1,t}^{(1)}|^2}{\|w_*\|^2}$ . Furthermore, by Lemma 4 and that  $|w_{\#1,t}^{(1)}| = c_{1,t}\|w_*\|^2$  and the definition of  $c_{2,t}$ ,

$\|w_{\#K,t}^{(k),\perp}\|_2 \lesssim \frac{|w_{\#1,t}^{(1),\perp}| \|w_{\#K,0}^{(k),\perp}\|}{|w_{\#K,0}^{(k),\perp}|} \frac{1}{\psi^t} = \frac{|w_{\#1,t}^{(1),\perp}| \|w_{\#K,0}^{(k),\perp}\|}{|w_{\#K,0}^{(k),\perp}| [1] \|w_*\|} \frac{1}{\psi^t} \leq c_{2,t} \|w_*\|$ . So we have that

$$\begin{aligned} \|W_t^{\#K}\|_F^2 &= \sum_{k=1}^K \|w_{\#K,t}^{(k)}\|^2 = \sum_{k=1}^K (\|w_{\#K,t}^{(k),\perp}\|^2 + \frac{1}{\|w_*\|^2} |w_{\#K,t}^{(k),\parallel}|^2) \\ &\lesssim \sum_{k=1}^K \left( \left( \frac{|w_{\#1,t}^{(1),\perp}| \|w_{\#K,0}^{(k),\perp}\|}{|w_{\#K,0}^{(k),\perp}| [1] \|w_*\|} \frac{1}{\psi^t} \right)^2 + \frac{1}{\|w_*\|^2} |w_{\#K,t}^{(k),\parallel}|^2 \right) \leq K \|w_*\|^2 c_{2,t}^2 + \sum_{k=1}^K \frac{|w_{\#K,t}^{(k),\parallel}|^2}{\|w_*\|^2} \\ &\lesssim K \|w_*\|^2 (c_{2,t}^2 + c_{1,t}^2), \end{aligned} \quad (17)$$

where the last inequality uses (11). Therefore, by combining (15,16,17),

$$\text{dist}^2(W_t^{\#1}, w_*) - \text{dist}^2(W_t^{\#K}, w_*) \gtrsim \|w_*\|^2 (2c_{1,t}(\sqrt{(1-2t\theta)}\sqrt{K}-1) - K(c_{1,t}^2 + c_{2,t}^2) + c_{1,t}^2). \quad (18)$$

■

#### 4.4. Is over-parametrization equivalent to using a larger step size?

Following the simulation as Figure 1, we tried different values of step sizes

$$\eta = \{1.0, 0.5, 0.4, 0.3, 0.2, 0.1, 0.05, 0.01, 0.005, 0.001\}$$

for each model with different number of neurons  $K = \{1, 3, 10\}$ . We report the quantity (9) over iteration  $t$ ,

$$\text{vec}(w_* q_t^\top - W_t^{\#K})^\top \nabla^2 f(W_t^{\#K}) \text{vec}(w_* q_t^\top - W_t^{\#K}),$$

where  $\nabla^2 f(W_t^{\#K}) \in \mathcal{R}^{dK \times dK}$  is the Hessian and  $w_* q_t^\top$  is the closet global optimal solution to  $W_t^{\#K}$  and the notation  $\text{vec}(\cdot)$  represents the vectorization operation of its matrix argument. Recall that the quantity can be viewed as a measure of the strong convexity as mentioned in the main text. Specifically, if the quantity is larger than 0, then it suggests that the current optimization landscape is strongly convex with respect to  $w_* q_t^\top$ .

We found out that gradient descent with step size  $\eta = \{1.0, 0.5, 0.4, 0.3\}$  either diverges or cannot converge towards zero testing error for all  $K = \{1, 3, 10\}$ , which means that  $\eta = 0.2$  is basically the best step size of gradient descent for each model  $K$ . So the result suggests that even under optimal tuning of the step size  $\eta$  for each model  $K$ , gradient descent for an over-parametrized model converges faster than the case for a smaller model. We show some results of using different  $\eta$  on Figure 6, Figure 7, and Figure 8. Based on the empirical results, we conclude that the acceleration due to over-parametrization cannot be simply explained by that gradient descent uses a larger *effective* step size, as the impacts due to parameters  $\eta$  and  $K$  seem to be complementary in the experiment.

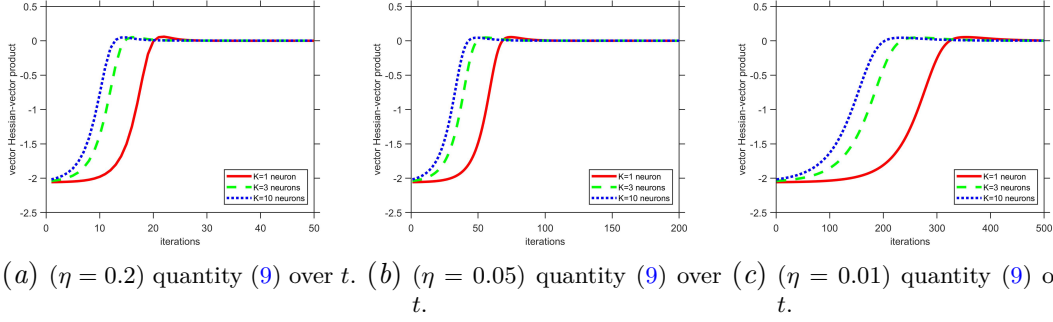


Figure 6: Gradient descent with different values of the step size. Note that the scales of the horizontal axes are different. Both the step size  $\eta$  and the degree of over-parametrization affect the time that gradient descent enters the linear convergence regime. A larger step size  $\eta$  and a larger number of neurons  $K$  help gradient descent to make progress faster.




Figure 7: Gradient descent with  $\eta = 0.5$ . We see that gradient descent cannot converge to zero training (and testing) error. That is, the step size is too large to converge to a global optimal solution. Interestingly, we still observe that gradient descent requires fewer iterations to get closer to a global optimal point for a larger model, though it does not converge to a global optimal point using the large step size.

## 5. Conclusion

We study over-parametrization for learning a single teacher neuron with quadratic activation. We answer why gradient descent can achieve a faster convergence for training a larger network, and we also show the acceleration due to over-parametrization cannot be simply explained by that gradient descent uses a larger effective step size. We hope our work can serve as a good starting point of understanding when and why over-parametrization leads to acceleration in modern non-convex optimization. Future works include considering different activation and multiple teacher neurons.




Figure 8: We plot the number of iterations required for the metric  $\text{vec}(w_* q_t^\top - W_t^{\#K})^\top \nabla^2 f(W_t^{\#K}) \text{vec}(w_* q_t^\top - W_t^{\#K})$  to be positive (the y-axis) under different values of the step size  $\eta$  (the x-axis).

## Acknowledgments

The authors acknowledge support of NSF IIS Award 1910077. The work was performed when JW was at Georgia Tech.

## References

- Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. *ICML*, 2018.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. *NerurIPS*, 2019.
- Alon Brutzkus and Amir Globerson. Why do larger models generalize better? a theoretical perspective via the xor problem. *ICML*, 2019.
- Emmanuel J. Candès, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 2013.
- Yuxin Chen, Yuejie Chi, Jianqing Fan, Cong Ma, and Yuling Yan. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 2019.
- Damek Davis, Dmitriy Drusvyatskiy, and Courtney Paquette. The nonsmooth landscape of phase retrieval. *IMA Journal on Numerical Analysis*, 2018.
- Simon Du and Jason Lee. On the power of over-parametrization in neural networks with quadratic activation. *ICML*, 2018.

- Matt Emschwiller, David Gamarnik, Eren C. Kızıldag, and Ilias Zadik. Neural networks and polynomial regression. demystifying the overparametrization phenomena. *arXiv:2003.10523*, 2020.
- Albert Fannjiang and Thomas Strohmer. The numerics of phase retrieval. *Acta Numerica*, 2020.
- Spencer Frei, Yuan Cao, and Quanquan Gu. Agnostic learning of a single neuron with gradient descent. *arXiv:2005.14426*, 2020.
- David Gamarnik, Eren C. Kızıldag, and Ilias Zadik. Stationary points of shallow neural networks with quadratic activation function. *arXiv:1912.01599*, 2019.
- Rong Ge, Runzhe Wang, and Haoyu Zhao. Mildly overparametrized neural nets can memorize training data efficiently. *arXiv:1909.11837*, 2019.
- Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. *NeurIPS*, 2019.
- Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the relu in polynomial time. *COLT*, 2017.
- Surbhi Goel, Adam Klivans, and Raghu Meka. Time/accuracy tradeoffs for learning a relu with respect to gaussian marginals. *NeurIPS*, 2019.
- Sebastian Goldt, Madhu Advani, Andrew M. Saxe, Florent Krzakala, and Lenka Zdeborova. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. *NeurIPS*, 2019.
- Suriya Gunasekar, Blake Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nathan Srebro. Implicit regularization in matrix factorization. *NIPS*, 2017.
- Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *NIPS*, 2017.
- Sham M. Kakade, Varun Kanade, Ohad Shamir, and Adam Kalai. Efficient learning of generalized linear and single index models with isotonic regression? *NeurIPS*, 2011.
- Seyed Mohammadreza Mousavi Kalan, Mahdi Soltanolkotabi, and A. Salman Avestimehr. Fitting relus via sgd and quantized sgd. *ISIT*, 2019.
- Abbas Kazemipour, Brett Larsen, and Shaul Druckmann. No spurious local minima in deep quadratic networks. *arXiv:2001.00098*, 2019.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- Yuanxin Li, Cong Ma, Yuxin Chen, and Yuejie Chi. Nonconvex matrix factorization from rank-one measurements. *AISTATS*, 2019.



- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. *COLT*, 2018.
- Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. *NeurIPS*, 2014.
- Stefano Sarao Mannelia, Eric Vanden-Eijnden, and Lenka Zdeborová. Optimization and generalization of shallow neural networks with quadratic activation functions. *arXiv:2006.15459*, 2020.
- Song Mei, Yu Bai, and Andrea Montanari. The landscape of empirical risk for non-convex losses. *PNAS*, 2018.
- Quynh Nguyen and Matthias Hein. The loss surface of deep and wide neural networks. *ICML*, 2017.
- Yoav Shechtman, Yonina C. Eldar, Oren Cohen, Henry Nicholas Chapman, Jianwei Miao, and Mordechai Segev. Phase retrieval with application to optical imaging: a contemporary overview. *IEEE signal processing magazine*, 2015.
- Mahdi Soltanolkotabi. Algorithms and theory for clustering and nonconvex quadratic programming. *Stanford University Ph. D. Dissertation*, 2014.
- Mahdi Soltanolkotabi. Learning relus via gradient descent. *NeurIPS*, 2017.
- Mahdi Soltanolkotabi, Adel Javanmard, and Jason D. Lee. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *IEEE Transactions on Information Theory*, 2018.
- Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *IEEE ISIT*, 2016.
- Yuandong Tian. Student specialization in deep rectified networks with finite width and input dimension. *ICML*, 2020.
- Luca Venturi, Afonso S Bandeira, and Joan Bruna. Spurious valleys in one-hidden-layer neural network optimization landscapes. *JMLR*, 2019.
- Chris D. White, Sujay Sanghavi, and Rachel Ward. The local convexity of solving systems of quadratic equations. *Results in Mathematics*, 2016.
- Gilad Yehudai and Ohad Shamir. Learning a single neuron with gradient methods. *arXiv:2001.05205*, 2020.