

Global Convergence and Geometric Characterization of Slow to Fast Weight Evolution in Neural Network Training for Classifying Linearly Non-Separable Data

Ziang Long^{1*} and Penghang Yin^{2†} and Jack Xin¹

Abstract

In this paper, we study the dynamics of gradient descent in learning neural networks for classification problems. Unlike in existing works, we consider the linearly non-separable case where the training data of different classes lie in orthogonal subspaces. We show that when the network has sufficient (but not exceedingly large) number of neurons, (1) the corresponding minimization problem has a desirable landscape where all critical points are global minima with perfect classification; (2) gradient descent is guaranteed to converge to the global minima in this case. Moreover, we discovered a geometric condition on the network weights so that when it is satisfied, the weight evolution transitions from a slow phase of weight direction spreading to a fast phase of weight convergence. The geometric condition says that the convex hull of the weights projected on the unit sphere contains the origin.

1 Introduction

Deep neural networks (DNN) have achieved remarkable performances in image and speech classification tasks among other AI applications in recent years; for examples, see Hochreiter and Schmidhuber (1997); Krizhevsky, Sutskever, and Hinton (2012); Ren et al. (2015); Silver et al. (2016). Although there have been numerous theoretical contributions to understand their success, the learning process in the actual network training remains largely empirical. One interesting phenomenon is that over-parametrized DNN's trained by stochastic gradient descent generalize Neyshabur et al. (2018); Nguyen, Mukkamala, and Hein (2018) instead of overfitting the training data contrary to conventional statistical learning. Though several convergence results are proved in the over-parameterized regime for deep networks (Du et al., 2018; Li and Liang, 2018; Allen-Zhu, Li, and Song, 2018), the network weights move only in a small neighborhood of the random initialization and so their dynamics are very localized. Partly, this may be attributed to the exceedingly large number of neurons in convergence theory, far surpassing what is used in practice where the weights evolve

significantly from random start through hundreds of epochs in training to reach best prediction accuracy.

Our work here addresses how the weights evolve towards a global minimum of loss function as the number of neurons increases from the feature dimension (the least necessary) to the over-parametrized regime. To facilitate analysis, our model network structure is motivated by (Brutzkus et al., 2018) on classifying linearly separable data. We instead study a linearly non-separable binary classification problem with an emphasis on the dynamics of weights in terms of the two time scales of evolution and a geometric characterization of the transition time. Our training data of the two classes will lie in orthogonal sub-spaces, which extends the data configuration in Brutzkus and Globerson (2018) where the subspace of each class is one dimensional for an XOR detection problem. Orthogonality of input data from the two classes implies that the training process in each class can be analyzed independently of the other. In the one-dimensional case Brutzkus and Globerson (2018), each weight update does not increase the loss on any sample point. In the multi-dimensional case here, we find that during gradient descent weight update, it is not possible that the loss is non-increasing in the point-wise sense (on each input data). Instead, the population loss is decreasing (i.e. in the sense of expectation). The population loss here is based on the hinge loss function and the network activation function is ReLU. Under a mild non-degenerate data condition, we prove that all critical points of our non-convex and non-smooth population loss function are global minima. Similar landscapes (a local minimum is a global minimum) are known for deep networks with activation functions that are either strictly convex (Liang et al., 2018), or real analytic and strictly increasing Nguyen, Mukkamala, and Hein (2018).

1.1 Prior Works and Our Contributions

In DNN training, one observes that the network learning consists of alternating phases: plateaus where the validation error remains fairly constant and periods of rapid improvement where a lot of progress is made over a few epochs. Prior to our work, (des Combes et al., 2019) studied slow and fast weight dynamics in a solvable model while minimizing a binary cross entropy or hinge loss function on linearly sep-

*University of California, Irvine

†University at Albany, SUNY

arable data. In the regression context, Brutzkus and Globerson (2017) came across such two time-scale phenomenon in training a two-linear-layer convolutional network with prescribed ground truth and unit Gaussian input data. This particular data assumption makes it possible to readily derive the closed-form expressions of the population loss and gradient, and then analyze the energy landscape and convergence of the gradient descent algorithm.

In this work, we study network weight dynamics in training a one-hidden-layer ReLU network via hinge loss minimization on binary classification of linearly non-separable data lying in two orthogonal sub-spaces. Our main contributions are:

- We discovered a geometric condition (GC) to characterize the transition time T from the first (slow) phase of weight evolution to the second fast weight convergence. The condition says that the convex hull of the weights on the unit sphere contains the origin, see Fig. 1 for an illustration. Equivalent geometric conditions are also derived (Lemma 1). In the first (slow) phase, the weight directions spread out over the unit sphere to satisfy GC.
- We obtain upper bound on T in terms of data distribution function provided that the network weights are uniformly bounded during training which we observed numerically.
- We give probabilistic bounds on the validity of geometric condition for random initialization, which suggests that the larger the number of neurons, the more likely GC holds and the earlier the fast phase of evolution begins.
- We prove the global convergence of gradient descent training algorithm under the uniformly bounded weight assumption. In case of positive network bias, we prove a global Lipschitz gradient property of the loss function and sub-sequential convergence of weights to a global minimum. In case of zero network bias, we prove that the loss function has Lipschitz gradient away from the origin and is piece-wise C^1 . Moreover, the training loss converges to zero at a rate faster than linear.
- We prove that all critical points of the population loss function are global minima under a non-degenerate data condition.
- We provide numerical examples to substantiate our theory and illustrate the weight dynamics as the network size increases towards the over-parametrized regime.

Organization. In section 2, we introduce the settings of the classification problem, including the assumptions on the data and network architecture. In section 3, we state the main results regarding the convergence guarantee of the gradient descent algorithm for training the neural net in the cases of with and without a bias term in the linear layer. In section 4, we present preliminaries about the landscape of the training loss function. The convergence analysis of main results will be sketched in section 5. In section 6, we substantiate our theoretical findings with numerical simulations. All the technical proofs are detailed in the appendix.

Notations. We denote by \mathcal{S}^{d-1} the unit sphere in \mathbb{R}^d , and $|\mathcal{S}^{d-1}|$ the area of the unit sphere in the corresponding dimension. For any finite dimensional linear space $V \subseteq \mathbb{R}^d$,

we define V^k to be the collection of matrices of form $[\mathbf{x}_1, \dots, \mathbf{x}_k] \in \mathbb{R}^{d \times k}$, where $\mathbf{x}_j \in V$ is the j -th column vector. For any set \mathcal{X} , $\mathbb{1}_{\mathcal{X}}(x) = 1$ if $x \in \mathcal{X}$ else 0, is the indicator function of \mathcal{X} . For any vector $\mathbf{x} \in \mathbb{R}^d$, we denote $|\mathbf{x}|$ be the ℓ_2 norm of \mathbf{x} . For a matrix $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k] \in \mathbb{R}^{d \times k}$, $|\mathbf{W}| := \sum_{j=1}^k |\mathbf{w}_j|$ is the column-wise ℓ_2 -norm sum.

2 Problem Setup

In this section, we consider the binary classification problem in the $2d$ -dimensional space $\mathcal{X} = \mathbb{R}^{2d}$. Let $\mathcal{Y} = \{\pm 1\}$ be the set of labels, and let \mathcal{D}_1 and \mathcal{D}_2 be two probabilistic distributions over $\mathcal{X} \times \mathcal{Y}$. Throughout this paper, we make the following assumptions on the data:

1. **(Separability)** There are two orthogonal d -dimensional subspaces $V_i \subseteq \mathcal{X}$ for $i \in [2]$ with $V_1 \oplus V_2 = \mathcal{X}$, such that

$$\mathbb{P}_{(\mathbf{x}, y) \sim \mathcal{D}_i} [\mathbf{x} \in V_i \text{ and } y = (-1)^{i+1}] = 1.$$

2. **(Boundedness of data)** There exist positive constants m and M , such that

$$\mathbb{P}_{(\mathbf{x}, y) \sim \mathcal{D}_i} [m \leq |\mathbf{x}| \leq M] = 1$$

for $i = 1, 2$.

3. **(Boundedness of p.d.f.)** For $i = 1, 2$, let p_i be the probability density function of distribution \mathcal{D}_i restricted on V_i . For any $\mathbf{x}_i \in V_i$ with $m < |\mathbf{x}_i| < M$, it holds that

$$0 < p_{\min} \leq p_i(\mathbf{x}_i) \leq p_{\max} < \infty.$$

Later on, we denote \mathcal{D} to be the evenly mixed distribution of \mathcal{D}_1 and \mathcal{D}_2 .

We consider a two-layer neural network with $2k$ hidden neurons. Denote by $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k, \mathbf{u}_1, \dots, \mathbf{u}_k] \in \mathbb{R}^{2d \times 2k}$ the weight matrix in the hidden layer. For any input data $\mathbf{x} \in \mathcal{X} = \mathbb{R}^{2d}$, the neural net outputs

$$f(\mathbf{W}; \mathbf{x}) = \sum_{j=1}^k \sigma(\langle \mathbf{w}_j, \mathbf{x} \rangle - b) - \sum_{j=1}^k \sigma(\langle \mathbf{u}_j, \mathbf{x} \rangle - b), \quad (1)$$

where $\sigma := \max(\cdot, 0)$ is the ReLU function acting element-wise, and the bias $b \geq 0$ is a constant. The prediction is given by the network output label

$$\hat{y}(\mathbf{W}, \mathbf{x}) = \text{sign}(f(\mathbf{W}; \mathbf{x})),$$

ideally $\hat{y}(x) = 1$ if $x \in V_1$, $\hat{y}(x) = -1$ if $x \in V_2$. The classification accuracy in percentage is the frequency that this occurs (when network output label \hat{y} matches the true label) on a validation data set. Given the data sample $\{\mathbf{x}, y\}$, the associated hinge loss function reads

$$l(\mathbf{W}; \{\mathbf{x}, y\}) := \max\{1 - y \cdot f(\mathbf{W}; \mathbf{x}), 0\}. \quad (2)$$

For network training, we consider the gradient descent algorithm with step size $\eta > 0$

$$\mathbf{W}^{(t)} = \mathbf{W}^{(t-1)} - \eta \nabla l(\mathbf{W}^{(t-1)}) \quad (3)$$

to solve following population loss minimization problem

$$\min_{\mathbf{W} \in \mathbb{R}^{2d \times 2k}} l(\mathbf{W}) := \mathbb{E}_{\{\mathbf{x}, y\} \sim \mathcal{D}} [l(\mathbf{W}; \{\mathbf{x}, y\})], \quad (4)$$

where the sample loss function $l(\mathbf{W}; \{\mathbf{x}, y\})$ is given by (2).

3 Main Results

Although (4) is a non-convex optimization problem, we show that under mild conditions, the gradient descent algorithm (3) converges to a global minimum with zero classification error. Specifically, we consider two different networks with a positive bias $b > 0$ (Theorem 1) and without a bias (Theorem 2), respectively. For both cases, we have the fact that any critical point of problem (4) is a global minimum (Proposition 1). The key difference between these two cases is that the population loss function has Lipschitz continuous gradient (Lemma 2) when $b > 0$, whereas this desirable property does not hold otherwise. For the latter case $b = 0$, we present a totally different analysis based on a geometric condition proved emergent during the training process (Proposition 3). Under this geometric condition, the objective value converges super-linearly to zero (Proposition 4).

Theorem 1. *Suppose the bias term $0 < b < \frac{1}{2kM}$ in (1), and $\{\mathbf{W}^{(t)}\}$ generated by the algorithm (3) are bounded uniformly in t . If there exist at least two data points $\mathbf{x}_i \in V_i$ and indices $1 \leq j_i \leq k$, for $i = 1, 2$, such that $|\langle \mathbf{w}_{j_1}^{(0)}, \mathbf{x}_1 \rangle| > b$ and $|\langle \mathbf{u}_{j_2}^{(0)}, \mathbf{x}_2 \rangle| > b$, then there exists some $\eta_0(k, b, p_{\max}, M) > 0$ such that for the learning rate $\eta < \eta_0$, $\{l(\mathbf{W}^{(t)})\}$ monotonically decreases to 0, and*

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\{\mathbf{x}, y\} \sim \mathcal{D}} [\hat{y}(\mathbf{W}^{(t)}, \mathbf{x}) \neq y] = 0.$$

Theorem 2. *Suppose the bias term in (1) satisfies $b = 0$, and $|\mathbf{W}^{(t)}| \leq R$ for all t . If there exist at least two data points $\mathbf{x}_i \in V_i$ and indices $1 \leq j_i \leq k$, such that $\langle \mathbf{w}_{j_1}^{(0)}, \mathbf{x}_1 \rangle \neq 0$ and $\langle \mathbf{u}_{j_2}^{(0)}, \mathbf{x}_2 \rangle \neq 0$ for $i = 1, 2$, then*

$$\sum_{t=1}^{\infty} l(\mathbf{W}^{(t)}) < C(d, M, R, p_{\min}, p_{\max}, \eta) \quad (5)$$

and

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\{\mathbf{x}, y\} \sim \mathcal{D}} [\hat{y}(\mathbf{W}, \mathbf{x}) \neq y] = 0.$$

Remark 1. *The assumptions on the initialization $|\langle \mathbf{w}_{j_1}^{(0)}, \mathbf{x}_1 \rangle| > b$ and $|\langle \mathbf{u}_{j_1}^{(0)}, \mathbf{x}_2 \rangle| > b$ in both theorems are natural. The assumptions guarantee that at least some \mathbf{w}_j neurons are activated by input data of type 1 and some \mathbf{u}_j neurons are activated by input data of type 2. Without them, the algorithm suffers zero gradient and fails to update.*

Remark 2. *Theorem 2 does not explicitly require the learning rate η to be small. However, a larger learning rate will implicitly result in a larger bound R for $|\mathbf{W}^{(t)}|$, which contributes to accumulative objective C in (5).*

4 Preliminaries

4.1 Decoupling

Let l_i be the population loss function of data type i with the label $y = (-1)^i$, $i = 1, 2$. More precisely,

$$l_i(\mathbf{W}) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_i} [\max\{1 + y \cdot f(\mathbf{W}, \mathbf{x}), 0\}].$$

Thus, we can rewrite the loss function as

$$l(\mathbf{W}) = \frac{1}{2} (l_1(\mathbf{W}) + l_2(\mathbf{W})).$$

Note that the population loss function

$$\begin{aligned} l_i(\mathbf{W}) &= \int_{V_i} l(\mathbf{W}, \mathbf{x}) p_i(\mathbf{x}) d\mathbf{x} \\ &= \int_{V_i \cap \{l(\mathbf{W}, \mathbf{x}) > 0\}} (1 - y \cdot f(\mathbf{W}, \mathbf{x})) p_i(\mathbf{x}) d\mathbf{x} \end{aligned}$$

has no closed-form solution even if p is a constant function on its support. We cannot use closed-form formula to analyze the learning process, which makes our work different from many other works.

Lemma 1. *For both $i = 1, 2$, if $\mathbf{W}_i^* \in \mathbb{R}^{2d \times 2k}$ solves the optimization problem*

$$\min_{\mathbf{W} \in V_i^{2k}} l_i(\mathbf{W}),$$

then $\mathbf{W}^* = \mathbf{W}_1^* + \mathbf{W}_2^*$ solves the original optimization problem

$$\min_{\mathbf{W} \in \mathbb{R}^{2d \times 2k}} l(\mathbf{W}).$$

Proof of Lemma 1. Note that $\mathcal{X} = \mathbb{R}^{2d} = V_1 \oplus V_2$, we decompose

$$\mathbf{w}_j = \mathbf{w}_{j,1} + \mathbf{w}_{j,2} \quad \text{and} \quad \mathbf{u}_j = \mathbf{u}_{j,1} + \mathbf{u}_{j,2}$$

where $\mathbf{w}_{j,1}, \mathbf{u}_{j,1} \in V_1$ and $\mathbf{w}_{j,2}, \mathbf{u}_{j,2} \in V_2$.

Since V_1 and V_2 are orthogonal spaces, we have

$$\langle \mathbf{w}_j, \mathbf{x} \rangle = \begin{cases} \langle \mathbf{w}_{j,1}, \mathbf{x} \rangle & \text{if } \mathbf{x} \in V_1 \\ \langle \mathbf{w}_{j,2}, \mathbf{x} \rangle & \text{if } \mathbf{x} \in V_2 \end{cases}$$

Similarly,

$$\langle \mathbf{u}_j, \mathbf{x} \rangle = \begin{cases} \langle \mathbf{u}_{j,1}, \mathbf{x} \rangle & \text{if } \mathbf{x} \in V_1 \\ \langle \mathbf{u}_{j,2}, \mathbf{x} \rangle & \text{if } \mathbf{x} \in V_2 \end{cases}$$

Letting $\mathbf{W}_i = [\mathbf{w}_{1,i}, \dots, \mathbf{w}_{k,i}, \mathbf{u}_{1,i}, \dots, \mathbf{u}_{k,i}]$, we have

$$f(\mathbf{W}; \mathbf{x}) = \begin{cases} f(\mathbf{W}_1, \mathbf{x}) & \text{if } \mathbf{x} \in V_1 \\ f(\mathbf{W}_2, \mathbf{x}) & \text{if } \mathbf{x} \in V_2 \end{cases}$$

Hence,

$$\begin{aligned} & \min_{\mathbf{W} \in \mathbb{R}^{2k \times 2d}} l(\mathbf{W}) \\ &= \frac{1}{2} \min_{\mathbf{W} \in \mathbb{R}^{2k \times 2d}} (l(\mathbf{W}_1) + l(\mathbf{W}_2)) \\ &= \frac{1}{2} \min_{\mathbf{W} \in \mathbb{R}^{2k \times 2d}} (l_1(\mathbf{W}_1) + l_2(\mathbf{W}_2)) \\ &= \frac{1}{2} \min_{\mathbf{W}_1 \in \mathbb{R}^{2k \times 2d}} l_1(\mathbf{W}_1) + \frac{1}{2} \min_{\mathbf{W}_2 \in \mathbb{R}^{2k \times 2d}} l_2(\mathbf{W}_2) \\ &= \frac{1}{2} \min_{\mathbf{W}_1 \in V_1^{2k}} l_1(\mathbf{W}_1) + \frac{1}{2} \min_{\mathbf{W}_2 \in V_2^{2k}} l_2(\mathbf{W}_2), \end{aligned}$$

which completes the proof. \square

The optimization problem can be decoupled into two independent problem of the same form. Therefore, it suffices to consider a simplified problem. Let $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k, \mathbf{u}_1, \dots, \mathbf{u}_k] \in V_1^{2k}$, where $\mathbf{w}_j, \mathbf{u}_j \in V_1 = \mathbb{R}^d$, the network output for the input data $\mathbf{x} \in V_1 = \mathbb{R}^d$ is given by

$$\tilde{f}(\mathbf{W}, \mathbf{x}) = \sum_{j=1}^k \sigma(\langle \mathbf{w}_j, \mathbf{x} \rangle - b) - \sum_{j=1}^k \sigma(\langle \mathbf{u}_j, \mathbf{x} \rangle - b). \quad (6)$$

Networks (1) and (6) are different, since the parameters in (1) are $\mathbf{W} \in \mathbb{R}^{2d \times 2k}$, whereas in (6), we have $\mathbf{W} \in \mathbb{R}^{d \times 2k}$. The corresponding input data are also in different dimensions. From now on, we just focus on the loss function associated with data of Class 1:

$$l_1(\mathbf{W}) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_1} \left[\max \left\{ 1 + \tilde{f}(\mathbf{W}, \mathbf{x}), 0 \right\} \right].$$

4.2 Landscape

The following Proposition 1 shows that while the loss function is non-convex, any critical points is in fact a global minimum, except for some degenerate cases.

Proposition 1. *Consider the neural network in (1). Assume $d > 1$, if \mathbf{W} is a critical point of $l(\mathbf{W})$ and there exists some $\mathbf{x}_i \in V_i$ such that $f(\mathbf{W}, \mathbf{x}_i) \neq 0$. Then, we have $l(\mathbf{W}) = 0$.*

Proof of Proposition 1. Assume \mathbf{W} is a critical point of $l(\mathbf{W})$, then for any $j \in [k]$, we have

$$\nabla_{\mathbf{w}_j} l(\mathbf{W}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\nabla_{\mathbf{w}_j} l(\mathbf{W}; \{\mathbf{x}, y\}) \right] = \mathbf{0}. \quad (7)$$

Let

$$\begin{aligned} \Omega_{\mathbf{v}} &= \{\mathbf{x} \in \mathcal{X} : \langle \mathbf{v}, \mathbf{x} \rangle > b\}, \\ \Omega_{\mathbf{W}} &= \{\mathbf{x} \in \mathcal{X} : y \cdot f(\mathbf{W}; \mathbf{x}) < 1\}, \end{aligned}$$

we have

$$\begin{aligned} \nabla_{\mathbf{w}_j} l(\mathbf{W}; \{\mathbf{x}, y\}) &= -y \cdot \mathbf{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \nabla_{\mathbf{w}_j} f(\mathbf{W}; \mathbf{x}) \\ &= -y \cdot \mathbf{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \mathbf{1}_{\Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{x}. \end{aligned} \quad (8)$$

Put equation (8) into equation (7), we have

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[y \cdot \mathbf{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \mathbf{1}_{\Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{x} \right] = \mathbf{0}.$$

Note that V_1 and V_2 are linearly independent, we have

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_i} \left[\mathbf{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \mathbf{1}_{\Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{x} \right] = \mathbf{0},$$

for $i = 1, 2$. Note that for any fixed \mathbf{v} , if $\mathbf{x}_i \in \Omega_{\mathbf{v}}$, then, we know $\langle \mathbf{v}, \mathbf{x}_i \rangle > b \geq 0$ so that all $\mathbf{x} \in \Omega_{\mathbf{v}}$ lie in same half space. Hence, we know

$$\mathbf{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \mathbf{1}_{\Omega_{\mathbf{w}_j}}(\mathbf{x}) = \mathbf{0},$$

for all $j \in [k]$ and almost sure all $\mathbf{x} \in \mathcal{X}$. The same arguments applies to \mathbf{u}_j gives

$$\mathbf{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \mathbf{1}_{\Omega_{\mathbf{u}_j}}(\mathbf{x}) = \mathbf{0},$$

for all $j \in [k]$ and almost sure all $\mathbf{x} \in \mathcal{X}$. Combine the above two equation, we know

$$\Omega_{\mathbf{W}} \cap \left[\bigcup_{j=1}^k (\Omega_{\mathbf{w}_j} \cup \Omega_{\mathbf{u}_j}) \right] = \Phi.$$

Recall the definitions of $\Omega_{\mathbf{W}}$ and $\Omega_{\mathbf{v}}$, we know $\mathbf{x} \in \Omega_{\mathbf{W}}$ if and only if $l(\mathbf{W}; \{\mathbf{x}, y\}) > 0$ and $\mathbf{x} \notin \bigcup_{j=1}^k (\Omega_{\mathbf{w}_j} \cup \Omega_{\mathbf{u}_j})$ implies $f(\mathbf{W}; \mathbf{x}) = 0$ so that $l(\mathbf{W}; \{\mathbf{x}, y\}) = 1$. Note that l is a continuous function, we know either $l(\mathbf{W}) = 0$ or $f(\mathbf{W}; \mathbf{x}) \equiv 0$. Now, we get the desired result. \square

The above Proposition is only effective when global minimum exists. The following proposition shows that the loss function has plenty of global minima.

Proposition 2. *Consider the network in (6). If the convex hull spanned by vertices \mathbf{w}_j 's contains a ball centered at the origin with radius $\frac{1+b}{m}$, and \mathbf{u}_j lies in a ball with radius $\frac{b}{M}$, then $l_1(\mathbf{W}) = 0$.*

The above proposition shows that if number of neurons is greater than the dimension of input data, then global minimum exists. Next, we study the smoothness of the loss function. The following proposition shows that as long as weights are bounded away from 0, then the loss function has Lipschitz gradient.

Lemma 2. *Consider the network in (1) with positive bias $0 < b < \frac{1}{2Mk}$. The loss function $l(\mathbf{W})$ is Lipschitz differentiable, i.e, there exists some constant $L(k, b, p_{\max}, M) > 0$ depending on b , such that*

$$|\nabla l(\mathbf{W}_1) - \nabla l(\mathbf{W}_2)| \leq L |\mathbf{W}_1 - \mathbf{W}_2|.$$

Note that Lipschitz differentiability in Lemma 2 does not hold for the case $b = 0$, as the gradient might be volatile near the origin.

5 Convergence Analysis for Non-Bias Case

With the Lipschitz differentiability shown in Lemma 2 in the case $b > 0$, it is not hard to prove the convergence result in Theorem 1. In this section, we focus on the non-bias case ($b = 0$) where the Lipschitz differentiability fails and sketch the convergence analysis.

Lemma 3. *Consider the network (6), $\{|\mathbf{w}_j^{(t)}|\}$ is non-decreasing. For any $r > 0$, choosing learning rate $\eta < \min \left\{ \frac{r}{C_p M^2}, \frac{r}{M} \right\}$, then $\{|\mathbf{u}_j^{(t)}|\}$ is non-increasing whenever $|\mathbf{u}_j| > r$, where C_p is a constant satisfying*

$$C_p = \max_{\mathbf{v} \in V_1, a \in \mathbb{R}} \int_{\{\langle \mathbf{v}, \mathbf{x} \rangle = a\}} p_1(\mathbf{x}) d\mathbf{x} = O(p_{\max} M^d).$$

The above theorem provides the dynamic information of the weights. When we input data with distribution \mathcal{D}_1 , the \mathbf{w}_j terms are becoming more useful for classification as the norm of \mathbf{w}_j grows larger on every iteration. On the other hand, \mathbf{u}_j terms serve only as noise. When \mathbf{u}_j are not small, the learning process guarantees the decreasing of population loss. When \mathbf{u}_j 's are small, they get trapped to a small




Figure 1: Geometric Condition in Lemma 4 (d=3)

region near the origin and contribute little to classification. One may wonder why we have to keep the \mathbf{u}_j terms, since they function only like noise. The reason that we could not drop the \mathbf{u}_j terms is that the roles of \mathbf{w}_j and \mathbf{u}_j terms switch with the input data coming from distribution \mathcal{D}_2 .

The learning process consists of two phases. In the beginning, since the weights are randomly distributed, there may well be some data that does not activate any neurons. The weights then automatically spread out. We call this process the first (slow learning) phase, during which the learning process is rather slow. The next Lemma lists four equivalent statements of a geometric condition. When the geometrical condition holds, we say that the learning process enters the second (fast learning) phase.

Lemma 4. Let $\tilde{\mathbf{w}}_j = \frac{\mathbf{w}_j}{|\mathbf{w}_j|} \in \mathcal{S}^{d-1}$ be k points in general position, where $1 \leq j \leq k$. Let Λ_W be the convex hull of $\{\tilde{\mathbf{w}}_j\}$ and $\{H_i\}_{i=1}^l$ be the facets of convex hull. The following statements are equivalent:

1. For any unit vector $\mathbf{n} \in \mathcal{S}^{d-1}$, there exist some j_1 and j_2 such that $\langle \mathbf{n}, \tilde{\mathbf{w}}_{j_1} \rangle > 0$ and $\langle \mathbf{n}, \tilde{\mathbf{w}}_{j_2} \rangle < 0$.
2. There exist no closed hemisphere that contains all $\tilde{\mathbf{w}}_j$.
3. $\mathbf{0}$ lies in the interior of Λ_W .
4. For each $0 \leq i \leq l$, we have $\mathbf{0}$ and all $\tilde{\mathbf{w}}_j$ lie on the same side of H_i .

Remark 3. If \mathbf{w}_j is initialized such that $\tilde{\mathbf{w}}_j$ is uniformly distributed on \mathcal{S}^{d-1} , then the probability that the geometric condition (GC) in Lemma 4 holds is

$$P_{\text{gc}} := \text{Prob}(\text{GC holds}) = 2^{1-k} \sum_{j=d}^{k-1} \binom{k-1}{j}.$$

In particular, at any fixed feature dimension d ,

$$\lim_{k \rightarrow \infty} P_{\text{gc}} = 1.$$

Proof of Remark 3. For any $J \in \{\pm 1\}^k$, we let $S_J(\{\tilde{\mathbf{w}}_j\}) = \{J_j \tilde{\mathbf{w}}_j : 1 \leq j \leq k\}$. From the choice of $\tilde{\mathbf{w}}_j$, we know all S_J have the same distribution, so that

$$\text{Prob}_{\{\tilde{\mathbf{w}}_j\}}(\text{GC holds}) = \text{Prob}_{S_J}(\text{GC holds}).$$

Hence, we can simplify the probability as follows

$$\text{Prob}(\text{GC holds}) = \frac{1}{2^n} \mathbb{E} \left[\sum_J \mathbb{1}_{\text{gc}}(S_J(\{\tilde{\mathbf{w}}_j\})) \right].$$

We claim that $\sum_J \mathbb{1}_{\text{gc}}(S_J(\{\tilde{\mathbf{w}}_j\}))$ is a constant independent of choice of $\tilde{\mathbf{w}}_j$ as long as they are in general position.

Let $\tilde{\mathbf{w}}_j \in \mathcal{S}^{d-1}$ where $1 \leq j \leq k$. Each $\tilde{\mathbf{w}}_j$ corresponds to a subspace H_j with codimension 1 in \mathbb{R}^d , that is

$$H_j = \{\mathbf{x} \in \mathbb{R}^d : \langle \tilde{\mathbf{w}}_j, \mathbf{x} \rangle = 0\}.$$

Note that any connected region of $(\cup_J H_J)^c$ corresponds to a choice of J such that the geometric condition fails. More precisely, for any given connected region D of $(\cup_J H_J)^c$, we know $\langle \tilde{\mathbf{w}}_j, \mathbf{x} \rangle$ is one sign for all $\mathbf{x} \in D$, so with

$$J_j = \text{sign}(\langle \tilde{\mathbf{w}}_j, \mathbf{x} \rangle),$$

we know \mathbf{x} correspond to $S_J(\{\tilde{\mathbf{w}}_j\})$ where the geometric condition fails. Hence, the number of connected regions of $(\cup_J H_J)^c$ is same as $2^n - \sum_J \mathbb{1}_{\text{gc}}(S_J(\{\tilde{\mathbf{w}}_j\}))$. By (Ho and Zimmerman, 2006), we have

$$\sum_J \mathbb{1}_{\text{gc}}(S_J(\{\tilde{\mathbf{w}}_j\})) = 2^n - 2 \sum_{j=0}^{d-1} \binom{k-1}{j} = 2 \sum_{j=d}^k \binom{k-1}{j}.$$

The desired result follows. \square

Per Remark 3, the more neurons the network has, higher possibility the geometric condition in Lemma 4 holds upon initialization. As a consequence, the learning process skips the first phase and goes straight to the fast learning phase. This explains why gradient descent for learning a over-parameterized network converges rapidly to a nearby critical point from random initialization.

The following proposition gives an upper bound on the maximum number of iterations for the learning process to enter the second phase.

Proposition 3. Let $b = 0$, and assume that $|\mathbf{W}^{(t)}| \leq R$ for all t . Let T_1 be the set of t such that $\{\mathbf{w}_j\}$ does not satisfy the geometric condition in Lemma 4, then $|T_1| \leq \frac{2RC_p}{\eta p_R}$, where p_R is a positive constant depending only on R and the p_{\min} . Also, the following estimate holds for p_R :

$$p_R = \Omega \left(\frac{p_{\min}}{\sqrt{d} (MR)^d} \right)$$

where C_p is the constant in Lemma 3. More precisely,

$$|T_1| < O \left(\frac{C_p \sqrt{d} R^{d+1} M^d}{\eta p_{\min}} \right).$$

At the beginning of the second phase, $\tilde{\mathbf{w}}_j$ are already evenly distributed. That is, the convex hull of $\tilde{\mathbf{w}}_j$ contains the origin, and any input data must at least activate some of the neurons. As long as an input data contributes to the loss, it also contributes to the gradient. So learning process becomes faster during the second phase. The following proposition shows the loss decays super-linearly.

Proposition 4. Let $b = 0$. Let T_2 be the set of t such that $\{\mathbf{w}_j\}$ satisfies the geometric condition in Lemma 4, and that $\|\mathbf{W}^t\|$ is upper-bounded by R at all t . Then:

$$\sum_{t \in T_2} l(\mathbf{W}^{(t)}) \leq 2\eta^{-1} R^2 M C_p.$$

6 Experiments

In this section, we report the results of our experiments on both synthetic data and MNIST data. On one hand, the experiments on synthetic data aim to show our assumptions are reasonable and our theoretical results coincide with simulations. On the other hand, the experiment on MNIST dataset exhibits an real world example of slow-to-fast training dynamics which suggests this phenomenon worth further study.

6.1 Synthetic Data

In light of Lemma 1, we simulate the training process of the simplified network (6) with $d = 2$. In our first and all simulations below, the training data set is given as

$$\hat{\mathcal{X}} = \left\{ \frac{20}{a} \left(\cos \frac{b\pi}{40}, \sin \frac{b\pi}{40} \right) : a \in [20] - [9], b \in [80] \right\}$$

with label $\hat{\mathcal{Y}} = \{1\}$. In our first simulation, we take entries of $\mathbf{W}^{(0)}$ to be standard normal i.e. $\mathbf{w}_j^{(0)}, \mathbf{u}_j^{(0)} \sim N(\mathbf{0}, I_2)$. We train the network (6) with gradient descent (3) in all our simulations, where we take the learning rate $\eta = \frac{1}{10|\hat{\mathcal{X}}|}$. Fig. 2 shows how many iterations algorithm (3) takes in searching for a global minima from the random initialization mentioned above. For each box, the red mark indicates the median, and the bottom and top of the box indicate the 25th and 75th percentiles, respectively. As we can see from the graph, as number of hidden neurons ($2k$) becomes larger, the algorithm (3) tends to need less iterations in searching for a global minima.




Figure 2: convergent iterations vs. number of neurons ($d = 2$)

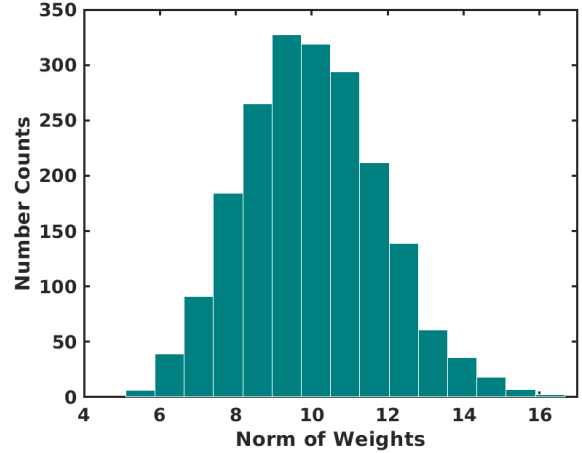


Figure 3: Right: histogram of norm of weights: $\max_t \|\mathbf{W}^{(t)}\|$ ($d = 2$ and $k = 4$).

In the second simulation, we compare the convergence speed with and without the geometric condition being satisfied. We introduce two initialization method: random initialization and half space initialization i.e. with $\hat{w}_{j,i}, \hat{u}_{j,i} \sim N(0, 1)$, random initialization takes $w_{j,i}^{(0)} = \hat{w}_{j,i}$ and $u_{j,i}^{(0)} = \hat{u}_{j,i}$ whereas half space initialization takes $w_{j,1}^{(0)} = |\hat{w}_{j,1}|$, $w_{j,2}^{(0)} = \hat{w}_{j,2}$, $u_{j,1}^{(0)} = |\hat{u}_{j,1}|$, and $u_{j,2}^{(0)} = \hat{u}_{j,2}$. We run the algorithm for 100 times with different numbers of hidden neurons using initialization methods, and report the means and standard variances of the number of iterations in Table 1. We see from Remark 3 how the P_{gc} increases when the number of hidden neurons grows. However, the half space initialization never satisfies the geometric condition, as all the weights lie in the same half space. A widely believed explanation on why a neural network can fit all training labels is that the neural network is over-parameterized. Our work explained one of the reasons why over-parameterization helps convergence: it helps the weights to spread more 'evenly' and quickly after initialization. Table 1 shows that when we randomly initialize, the iterations for convergence in gradient descent (3) come down a lot as the number of hidden neurons increases; much less so in half space initialization.

Our third simulation take specifically $2k = 8$. With 2000 runs we did a histogram of the maximum norm of \mathbf{W} during the training process shown in Fig. 3. In fact, our third simulation suggests our boundedness assumption on \mathbf{W} in Theorem 1 and Theorem 2 are reasonable.

In our last simulation, we take $k = 3$ so that there are in total 6 hidden neurons. We plot \hat{w}_j 's and \mathbf{u}_j 's in Fig. 4, where we plot \hat{w}_j 's instead of w_j 's since some of $|w_j|$'s are greater than one. Before algorithm (3) starts, the parameters in neural network (6) are initialized to be

$$\mathbf{w}_j^{(0)} = \mathbf{u}_j^{(0)} = \frac{3}{4} \left(\cos \frac{(2-j)\pi}{6}, \sin \frac{(2-j)\pi}{6} \right)$$

for $j \in \{1, 2, 3\}$. In Fig. 4, the tiny blue points are input

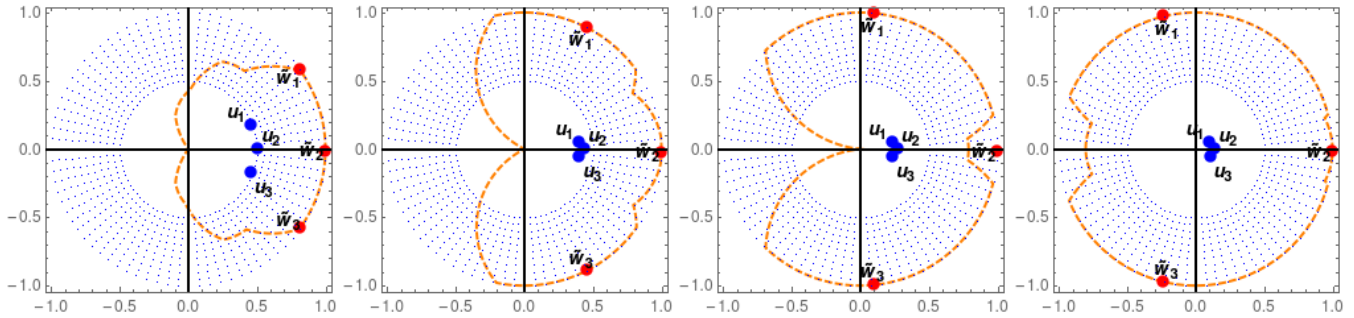


Figure 4: Dynamic of weights: \tilde{w}_j and u_j

Table 1: Iterations taken (mean \pm std) to convergence with random and half space initializations.

| # of Neurons ($2k$) | Random Init. | Half Space Init. |
|-----------------------|---------------------|---------------------|
| 6 | 578.90 \pm 205.43 | 672.41 \pm 226.53 |
| 8 | 423.96 \pm 190.91 | 582.16 \pm 200.81 |
| 10 | 313.29 \pm 178.67 | 550.19 \pm 180.59 |
| 12 | 242.72 \pm 178.94 | 517.26 \pm 172.46 |
| 14 | 183.53 \pm 108.60 | 500.42 \pm 215.87 |
| 16 | 141.00 \pm 80.66 | 487.42 \pm 220.48 |
| 18 | 126.52 \pm 62.07 | 478.25 \pm 202.71 |
| 20 | 102.09 \pm 32.32 | 412.46 \pm 195.92 |
| 22 | 90.65 \pm 28.01 | 454.08 \pm 203.00 |
| 24 | 82.93 \pm 26.76 | 416.82 \pm 216.58 |

data under Kelvin transformation: $\mathbf{x} \rightarrow \mathbf{x}^* = \frac{\mathbf{x}}{|\mathbf{x}|^2}$. Take $\mathbf{x} = \frac{1}{r}(\cos \theta, \sin \theta)$ so that under Kelvin transformation $\mathbf{x}^* = r(\cos \theta, \sin \theta)$. For convenience, we let $\tilde{\mathbf{x}} = r\mathbf{x} = (\cos \theta, \sin \theta)$. The orange dashed curve has expression in polar coordinates:

$$\rho(\theta) = \min \left\{ 1, \sigma \left(\tilde{f}(\mathbf{W}, \tilde{\mathbf{x}}) \right) \right\}.$$

Note that we are taking Hinge loss $l(\mathbf{W}, \{\mathbf{x}, 1\}) = 0$ if and only if $\tilde{f}(\mathbf{W}, \mathbf{x}) \geq 1$, i.e.

$$\tilde{f}(\mathbf{W}, \tilde{\mathbf{x}}) = r\tilde{f}(\mathbf{W}, \mathbf{x}) \geq r = |\mathbf{x}^*|.$$

Here, in our data set $\hat{\mathcal{X}}$, all data point have norm less than one under Kelvin transformation, so $l(\mathbf{W}, \{\mathbf{x}, 1\}) = 0$ if and only if $\rho(\theta) \geq |\mathbf{x}^*|$. This means, the blue points when surrounded by the orange dashed curve provide zero loss. In particular, when $\rho(\theta) = 1$, the population loss is 0.

6.2 MNIST Experiment

The two-phase dynamics we proved in our model does appear in deep network training on real (non-synthetic) data sets. In experiments on LeNet-5 and real linearly non-separable data set MNIST, we train LeNet-5 with SGD using a constant learning rate = 0.01 and without momentum or regularization. We show in Fig. 6 the loss value vs. iterations during training. At the early stage, the loss decays slowly, then the fast phase sets in after 400 iterations. Fig.




Figure 5: Slow-to-Fast transition on MNIST dataset

6 clearly supports our theory on the two-phase dynamics of gradient descent.

7 Summary

The slow and fast dynamics of neural network weights under gradient descent is critical for understanding the learning process. We performed the first theoretical study on training neural networks to classify linearly un-separable data sets away from the over-parametrized regime. We discovered a two time-scale phenomenon of network weights during gradient descent training: a slow phase where the weights spread out to satisfy a geometric condition, and a subsequent fast phase where the weights converge to a global minimum.

One direction for future work is to provide a concrete relation between the number of weights and the rate of convergence, and quantify the effect of over-parameterization on the rate of convergence. Another direction is to show that similar results hold on a more sophisticated linearly non-separable data set with a deeper neural network.

8 Acknowledgement

This work was partially supported by NSF grants IIS-1632935, DMS-1854434, DMS-1924548, and DMS-1924935.

References

- Allen-Zhu, Z.; Li, Y.; and Song, Z. 2018. A convergence theory for deep learning via over-parameterization. *CoRR* abs/1811.03962.
- Brutzkus, A., and Globerson, A. 2017. Globally optimal gradient descent for a convnet with gaussian inputs. *arXiv preprint arXiv:1702.07966*.
- Brutzkus, A., and Globerson, A. 2018. Over-parameterization improves generalization in the XOR detection problem. *CoRR* abs/1810.03037.
- Brutzkus, A.; Globerson, A.; Malach, E.; and Shalev-Shwartz, S. 2018. SGD learns over-parameterized networks that provably generalize on linearly separable data. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- des Combes, R. T.; Pezeshki, M.; Shabaniyan, S.; Courville, A.; and Bengio, Y. 2019. Convergence properties of deep neural networks on separable data.
- Du, S. S.; Zhai, X.; Póczos, B.; and Singh, A. 2018. Gradient descent provably optimizes over-parameterized neural networks. *CoRR* abs/1810.02054.
- Ho, C., and Zimmerman, S. 2006. On the number of regions in an m-dimensional space cut by n hyperplanes. *The Australian Mathematical Society Gazette* 33.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 1097–1105.
- Li, Y., and Liang, Y. 2018. Learning overparameterized neural networks via stochastic gradient descent on structured data. *CoRR* abs/1808.01204.
- Liang, S.; Sun, R.; Li, Y.; and Srikant, R. 2018. Understanding the loss surface of neural networks for binary classification. *CoRR* abs/1803.00909.
- Neysshabur, B.; Li, Z.; Bhojanapalli, S.; LeCun, Y.; and Srebro, N. 2018. Towards understanding the role of over-parametrization in generalization of neural networks. *CoRR* abs/1805.12076.
- Nguyen, Q.; Mukkamala, M. C.; and Hein, M. 2018. On the loss landscape of a class of deep neural networks with no bad local valleys. *CoRR* abs/1809.10749.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing systems*, 91–99.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484.

Appendix

We shall use the following notations in the following arguments. Let

$$\begin{aligned}\Omega_{\mathbf{W}} &= \{\mathbf{x} \in \mathcal{X} : l(\mathbf{W}, \mathbf{x}) > 0\}, \\ \Omega_{\mathbf{v}} &= \{\mathbf{x} \in \mathcal{X} : \langle \mathbf{v}, \mathbf{x} \rangle > b\} \\ \Omega_{\mathbf{W}}^j &= \Omega_{\mathbf{W}} \cap \Omega_{\mathbf{w}_j}, \\ \Omega_{\mathbf{W}}^{k+j} &= \Omega_{\mathbf{W}} \cap \Omega_{\mathbf{u}_j}.\end{aligned}$$

Proof of Theorem 1. By Lemma 1, we only need to prove the convergence of the simplified network (6). From Lemma 2, we know $l(\mathbf{W})$ has Lipschitz gradient. We can assume for any $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{2d \times 2k}$, we have

$$|\nabla l(\mathbf{W}_1) - \nabla l(\mathbf{W}_2)| \leq L |\mathbf{W}_1 - \mathbf{W}_2|.$$

As long as we take $\eta < \frac{L}{2}$ in algorithm (3), we know

$$\begin{aligned}l(\mathbf{W}^{(t+1)}) \\ \leq l(\mathbf{W}^{(t)}) - \left(\eta - \frac{\eta^2 L}{2}\right) |\nabla l(\mathbf{W}^{(t)})|^2 \\ \leq l(\mathbf{W}^{(t)}).\end{aligned}\quad (9)$$

Hence, $l(\mathbf{W}^{(t)})$ is monotonically decreasing. Therefore, for any convergent subsequence $\{\mathbf{W}^{(t_k)}\}$ with the limit \mathbf{W}_0 , there exists $l_0 \geq 0$ such that

$$\lim_{t \rightarrow \infty} l(\mathbf{W}^{(t)}) = \lim_{k \rightarrow \infty} l(\mathbf{W}^{(t_k)}) = l_0.$$

Now, we can take subsequence and limit on both side of equation (9), we get

$$l_0 \leq l_0 - \left(\eta - \frac{\eta^2 L}{2}\right) |\nabla l(\mathbf{W}_0)|^2.$$

By Lemma 3, $|\mathbf{w}_j|$ is increasing, whereas $|\mathbf{u}_j|$ is decreasing and got trapped in a small neighborhood of the origin. We know as long as $|\mathbf{w}_j^{(0)}| \neq 0$ for some j , then $f(\mathbf{W}_0, \cdot) \neq 0$. Hence, \mathbf{W}_0 is a critical point of $l(\mathbf{W})$. By Proposition 1, we know \mathbf{W}_0 is a global min. From inequality (9), we know the population loss is decreasing, so the entire sequence converge to zero loss:

$$\lim_{t \rightarrow \infty} l(\mathbf{W}^{(t)}) = 0.$$

Note that if for some $\mathbf{x} \in \mathcal{X}$ the classification is wrong, then $y \cdot f(\mathbf{W}^{(t)}, \{\mathbf{x}, y\}) \leq 0$, which means $l(\mathbf{W}^{(t)}, \{\mathbf{x}, y\}) \geq 1$. Since loss function is non-negative, by Markov's inequality, we have

$$\begin{aligned}& \mathbb{P}_{\{\mathbf{x}, y\} \sim \mathcal{D}} [\hat{y}(\mathbf{W}, \mathbf{x}) \neq y] \\ &= \mathbb{P} [y \cdot f(\mathbf{W}^{(t)}, \{\mathbf{x}, y\}) \leq 0] \\ &= \mathbb{P} [l(\mathbf{W}^{(t)}, \{\mathbf{x}, y\}) \geq 1] \\ &\leq l(\mathbf{W}^{(t)}).\end{aligned}$$

Now, we get the desired result. \square

Proof of Theorem 2. By Proposition 3, we know the iterates $\{\mathbf{W}^{(t)}\}$ stay in the first phase is bounded by $\frac{2RC_p}{\eta p_R}$. Also we know since weights are bounded, population loss is also bounded by RM . Combining Proposition 4 which shows the summation of loss values in phase two is bounded, the summation of all loss values in the learning process is bounded

$$\begin{aligned}& \sum_{t=0}^{\infty} l(\mathbf{W}^{(t)}) \\ & \leq \frac{2RC_p}{\eta p_R} \cdot RM + 2\eta^{-1} R^2 M(d+1)C_p \\ & < \infty.\end{aligned}$$

So,

$$\lim_{t \rightarrow \infty} l(\mathbf{W}^{(t)}) = 0.$$

Using similar arguments in proof of Theorem 1 completes the proof. \square

Proof of Proposition 2. If the convex hull spanned by vertices \mathbf{w}_j 's contains a ball centered at the origin with radius $\frac{1+b}{m}$, then for any $\mathbf{x} \in V_1$ with $|\mathbf{x}| \geq m$, there must exist some $1 \leq j_0 \leq k$ such that $\langle \mathbf{w}_{j_0}, \mathbf{x} \rangle \geq 1+b$ and thus $\sigma(\langle \mathbf{w}_{j_0}, \mathbf{x} \rangle - b) \geq 1$. Also, since all \mathbf{u}_j 's lie in a ball with radius $\frac{b}{M}$, for any $\mathbf{x} \in V_1$ and $|\mathbf{x}| \leq M$ we have $\sigma(\langle \mathbf{u}_j, \mathbf{x} \rangle - b) = 0$. Combine the results above, we know $\hat{f}(\mathbf{W}, \mathbf{x}) \geq 1$ for all $\mathbf{x} \in V_1$ and $m < |\mathbf{x}| < M$. \square

Proof of Lemma 2. We first calculate the true gradient

$$\begin{aligned}& \nabla_{\mathbf{w}_j} f(\mathbf{W}, \mathbf{x}) \\ &= \nabla_{\mathbf{w}_j} \sigma(\langle \mathbf{w}_j, \mathbf{x} \rangle - b) \\ &= \sigma'(\langle \mathbf{w}_j, \mathbf{x} \rangle - b) \mathbf{x} \\ &= \mathbb{1}_{\Omega_{\mathbf{W}}^j}(\mathbf{x}) \mathbf{x}.\end{aligned}$$

Recall the definition of $\Omega_{\mathbf{W}}$ and $\Omega_{\mathbf{v}}$ in proof of Proposition 1, we have

$$\begin{aligned}& \nabla_{\mathbf{w}_j} l(\mathbf{W}) \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\nabla_{\mathbf{w}_j} \max\{1 - y \cdot f(\mathbf{W}, \mathbf{x}), 0\}] \\ &= - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [y \cdot \mathbb{1}_{\Omega_{\mathbf{W}}}(\mathbf{x}) \cdot \nabla_{\mathbf{w}_j} f(\mathbf{W}, \mathbf{x})] \\ &= - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [y \cdot \mathbb{1}_{\Omega_{\mathbf{W}} \cap \Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{x}] \\ &= - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_1} [\mathbb{1}_{\Omega_{\mathbf{W}} \cap \Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{x}] + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_2} [\mathbb{1}_{\Omega_{\mathbf{W}} \cap \Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{x}].\end{aligned}$$

For $i = 1, 2$, let $\mathbf{W}_i = [\mathbf{w}_1^i, \dots, \mathbf{w}_1^i, \mathbf{u}_1^i, \dots, \mathbf{u}_k^i]$ so that $|\mathbf{W}_1 - \mathbf{W}_2| \leq \epsilon$, we claim there exists a constant L such that

$$|\nabla_{\mathbf{w}_j} l(\mathbf{W}_1) - \nabla_{\mathbf{w}_j} l(\mathbf{W}_2)| \leq L\epsilon.$$

First, we compare the gradient of two network with dif-

ferent weights \mathbf{W}_1 and \mathbf{W}_2

$$\begin{aligned}
& \left| \nabla_{\mathbf{w}_j} l(\mathbf{W}_1) - \nabla_{\mathbf{w}_j} l(\mathbf{W}_2) \right| \\
& \leq \left| \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_1} \left[\mathbb{1}_{\Omega_{\mathbf{W}_1} \cap \Omega_{\mathbf{w}_j^1}}(\mathbf{x}) \mathbf{x} \right] \right. \\
& \quad \left. - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_1} \left[\mathbb{1}_{\Omega_{\mathbf{W}_2} \cap \Omega_{\mathbf{w}_j^2}}(\mathbf{x}) \mathbf{x} \right] \right| \\
& + \left| \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_2} \left[\mathbb{1}_{\Omega_{\mathbf{W}_1} \cap \Omega_{\mathbf{w}_j^1}}(\mathbf{x}) \mathbf{x} \right] \right. \\
& \quad \left. - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_2} \left[\mathbb{1}_{\Omega_{\mathbf{W}_2} \cap \Omega_{\mathbf{w}_j^2}}(\mathbf{x}) \mathbf{x} \right] \right| \\
& \leq \underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_1} \left[\mathbb{1}_{\Omega_{\mathbf{W}_1} \Delta \Omega_{\mathbf{W}_2}}(\mathbf{x}) |\mathbf{x}| \right]}_{\textcircled{1}} \\
& + \underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_1} \left[\mathbb{1}_{\Omega_{\mathbf{w}_j^1} \Delta \Omega_{\mathbf{w}_j^2}}(\mathbf{x}) |\mathbf{x}| \right]}_{\textcircled{2}} \\
& + \underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_2} \left[\mathbb{1}_{\Omega_{\mathbf{W}_1} \Delta \Omega_{\mathbf{W}_2}}(\mathbf{x}) |\mathbf{x}| \right]}_{\textcircled{3}} \\
& + \underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_2} \left[\mathbb{1}_{\Omega_{\mathbf{w}_j^1} \Delta \Omega_{\mathbf{w}_j^2}}(\mathbf{x}) |\mathbf{x}| \right]}_{\textcircled{4}}
\end{aligned}$$

Next, we estimate $\textcircled{1}$ and $\textcircled{2}$. The difference between network output reads

$$\begin{aligned}
& |f(\mathbf{W}_1; \mathbf{x}) - f(\mathbf{W}_2; \mathbf{x})| \\
& = \left| \sum_{j=1}^k \sigma(\langle \mathbf{w}_j^1, \mathbf{x} \rangle - b) - \sigma(\langle \mathbf{w}_j^2, \mathbf{x} \rangle - b) \right. \\
& \quad \left. - \sigma(\langle \mathbf{u}_j^1, \mathbf{x} \rangle - b) + \sigma(\langle \mathbf{u}_j^2, \mathbf{x} \rangle - b) \right| \\
& \leq |\mathbf{W}_1^T \mathbf{x} - \mathbf{W}_2^T \mathbf{x}| \leq \epsilon |\mathbf{x}|
\end{aligned}$$

Note that $\mathbf{x} \in \Omega_{\mathbf{W}_1} \Delta \Omega_{\mathbf{W}_2}$ implies 1 lies between $f(\mathbf{W}_1; \mathbf{x})$ and $f(\mathbf{W}_2; \mathbf{x})$ so that

$$1 - \epsilon |\mathbf{x}| \leq f(\mathbf{W}_1; \mathbf{x}) \leq 1 + \epsilon |\mathbf{x}|,$$

W.l.o.g, we can assume $\epsilon > \frac{1}{2}$, let $\mathbf{x} = r\boldsymbol{\omega}$, where $r = |\mathbf{x}|$. Now,

$$\begin{aligned}
& \frac{d}{dr} f(\mathbf{W}; \mathbf{x}) \\
& = \left[\sum_{j=1}^k \mathbb{1}_{\Omega_{\mathbf{w}_j}}(\mathbf{x}) \mathbf{w}_j - \sum_{j=1}^k \mathbb{1}_{\Omega_{\mathbf{u}_j}}(\mathbf{x}) \mathbf{u}_j \right] \boldsymbol{\omega} \\
& \geq f(\mathbf{W}; \mathbf{x})/r - (k-2)b \\
& \geq \frac{1}{2M} - (k-2)b =: C > 0
\end{aligned}$$

So,

$$\begin{aligned}
\textcircled{1} & \leq \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbb{1}_{1-\epsilon|\mathbf{x}| < f(\mathbf{W}_1, \mathbf{x}) < 1+\epsilon|\mathbf{x}|}(\mathbf{x}) |\mathbf{x}| \right] \\
& \leq \int_{1-\epsilon|\mathbf{x}| < f(\mathbf{W}_1, \mathbf{x}) < 1+\epsilon|\mathbf{x}|} |\mathbf{x}| p_{\max} d\mathbf{x} \\
& = p_{\max} \int_{S^{d-1}} \frac{2\epsilon r^d}{C} dr d\boldsymbol{\omega} \\
& \leq 2 \left(|\mathcal{S}^{d-1}| \frac{M^d p_{\max}}{C} \right) \epsilon
\end{aligned}$$

For $\textcircled{2}$, w.l.o.g we can assume $|\mathbf{w}_j^1| \geq |\mathbf{w}_j^2|$. Note that when $|\mathbf{w}_j^1| \leq \frac{b}{M}$ then $\langle \mathbf{w}_j, \mathbf{x} \rangle - b \leq 0$ so that $\nabla_{\mathbf{w}_j} l(W) = \mathbf{0}$ and thus $\textcircled{2} = 0$. Hence, we only need to take care of the case when $|\mathbf{w}_j^1| \geq \frac{b}{M}$.

Note that $|\mathbf{w}_j^1 - \mathbf{w}_j^2| \leq \epsilon$ we know

$$\sin \theta \leq \frac{\epsilon M}{b},$$

where θ denotes the angle between \mathbf{w}_j^1 and \mathbf{w}_j^2 . Now, we have the following estimate

$$\textcircled{2} \leq p_{\max} \frac{\epsilon M}{b} |\mathcal{S}^{d-1}|.$$

For $\textcircled{3}$ and $\textcircled{4}$, we can use the same method. Together $\textcircled{1}$, $\textcircled{2}$, $\textcircled{3}$, and $\textcircled{4}$, we have the Lipschitz condition. \square

Proof of Lemma 3. We first define

$$C_p = \max_{v \in V_1, a \in \mathbb{R}} \int_{\langle v, \mathbf{x} \rangle = a} p_1(\mathbf{x}) dS \leq M^{d-1} p_{\max}.$$

Recall the definition of $\Omega_{\mathbf{W}}^j$, for almost all $\mathbf{x} \in \Omega_{\mathbf{W}}^j \cap V_1$, we have $\langle \tilde{\mathbf{w}}_j, \mathbf{x} \rangle > 0$.

$$\langle \tilde{\mathbf{w}}_j, \nabla_{\mathbf{w}_j} l_1(\mathbf{W}) \rangle = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_1} \left[\mathbb{1}_{\Omega_{\mathbf{W}}^j}(\mathbf{x}) \langle \tilde{\mathbf{w}}_j, \mathbf{x} \rangle \right] \leq 0.$$

Note we have from (3) that

$$\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} - \eta \nabla_{\mathbf{w}_j} l_1(\mathbf{W}^{(t)}).$$

So,

$$\begin{aligned}
|\mathbf{w}_j^{(t+1)}| & = \left\langle \mathbf{w}_j^{(t+1)}, \tilde{\mathbf{w}}_j^{(t+1)} \right\rangle \\
& \geq \left\langle \mathbf{w}_j^{(t+1)}, \tilde{\mathbf{w}}_j^{(t)} \right\rangle \\
& \geq \left\langle \mathbf{w}_j^{(t)}, \tilde{\mathbf{w}}_j^{(t)} \right\rangle \\
& = |\mathbf{w}_j^{(t)}|.
\end{aligned}$$

For \mathbf{u}_j , we know

$$|\nabla_{\mathbf{u}_j} l_1(\mathbf{W})| = \left| \mathbb{E} \left[\mathbb{1}_{\Omega_{\mathbf{W}}^{k+j}}(\mathbf{x}) \mathbf{x} \right] \right| \leq M \mathbb{P}[\Omega_{\mathbf{W}}^{k+j}],$$

where we omit the distribution $\mathbf{x} \sim \mathcal{D}_1$.

On the other hand, by definition of C_p , we know

$$\begin{aligned}
& |\langle \nabla_{\mathbf{u}_j} l_1(\mathbf{W}), \tilde{\mathbf{u}}_j \rangle| \\
&= \left| \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_1} \left[\mathbb{1}_{\Omega_{\mathbf{W}}^{k+j}}(\mathbf{x}) \langle \tilde{\mathbf{u}}_j, \mathbf{x} \rangle \right] \right| \\
&= \left| \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_1} \left[\int_0^\infty \mathbb{1}_{\{\lambda < \langle \tilde{\mathbf{u}}_j, \mathbf{x} \rangle\}}(\mathbf{x}) d\lambda \right] \right| \\
&= \left| \int_0^\infty \mathbb{P} \left[\{\mathbf{x} : \langle \tilde{\mathbf{u}}_j, \mathbf{x} \rangle < \lambda\} \cap \Omega_{\mathbf{W}}^{k+j} \right] d\lambda \right| \\
&\geq \frac{\mathbb{P}[\Omega_{\mathbf{W}}^{k+j}]}{2C_p}.
\end{aligned}$$

When $\eta < \frac{r}{M}$, we have

$$\begin{aligned}
& \langle \mathbf{u}_j - \eta \nabla_{\mathbf{u}_j} l(\mathbf{W}), \tilde{\mathbf{u}}_j \rangle \\
&= \langle \mathbf{u}_j, \tilde{\mathbf{u}}_j \rangle - \eta \langle \nabla_{\mathbf{u}_j} l(\mathbf{W}), \tilde{\mathbf{u}}_j \rangle \\
&> r - \eta M \mathbb{P} \left[\Omega_{\mathbf{W}}^{k+j} \right] > 0.
\end{aligned}$$

Now, we decompose $\nabla_{\mathbf{u}_j} l(\mathbf{W}^{(t)})$ into two parts,

$$\begin{aligned}
\nabla_{\mathbf{u}_j} l(\mathbf{W}^{(t)}) &= \underbrace{\langle \tilde{\mathbf{u}}_j^{(t)}, \nabla_{\mathbf{u}_j} l(\mathbf{W}^{(t)}) \rangle}_{\mathbf{n}} \tilde{\mathbf{u}}_j^{(t)} \\
&+ \underbrace{\left(\nabla_{\mathbf{u}_j} l(\mathbf{W}^{(t)}) - \langle \tilde{\mathbf{u}}_j^{(t)}, \nabla_{\mathbf{u}_j} l(\mathbf{W}^{(t)}) \rangle \tilde{\mathbf{u}}_j^{(t)} \right)}_{\boldsymbol{\nu}}.
\end{aligned}$$

So that when $\eta < \frac{r}{M^2 C_p}$, we have

$$\begin{aligned}
& \left| \mathbf{u}_j^{(t+1)} \right|^2 \\
&= \left| \mathbf{u}_j^{(t)} - \eta \nabla_{\mathbf{u}_j} l(\mathbf{W}^{(t)}) \right|^2 \\
&= \left| \mathbf{u}_j^{(t)} - \eta \mathbf{n} \right|^2 + |\eta \boldsymbol{\nu}|^2 \\
&= \left| \mathbf{u}_j^{(t)} \right|^2 - \eta \left(2 \left| \mathbf{u}_j^{(t)} \right| |\mathbf{n}| - \eta \left(|\mathbf{n}|^2 + |\boldsymbol{\nu}|^2 \right) \right) \\
&\leq \left| \mathbf{u}_j^{(t)} \right|^2 - \eta \left(2 \left| \mathbf{u}_j^{(t)} \right| \frac{\mathbb{P} \left[\Omega_{\mathbf{W}}^{k+j} \right]}{2C_p} - \eta M^2 \mathbb{P} \left[\Omega_{\mathbf{W}}^{k+j} \right]^2 \right) \\
&\leq \left| \mathbf{u}_j^{(t)} \right|^2,
\end{aligned}$$

as long as $|\mathbf{u}_j| \geq r$. Now, we get the desired result. \square

Proof of Lemma 4. (1 \Leftrightarrow 2) is trivial.

(1 \Rightarrow 3) Proof by contradiction. Assume $\mathbf{0} \notin \Lambda_{\mathbf{W}}^\circ$. Since $\Lambda_{\mathbf{W}}^\circ$ is an open convex set, and $\{\mathbf{0}\}$ is a convex set, we know from geometric form of Hahn-Banach Theorem, that there exist a closed hyper-plane that separates $\{\mathbf{0}\}$ and $\Lambda_{\mathbf{W}}^\circ$. Hence, there exist a unit vector $\mathbf{n} \in \mathcal{S}^{d-1}$, such that $\langle \mathbf{n}, \tilde{\mathbf{w}}_j \rangle \geq \langle \mathbf{n}, \mathbf{0} \rangle = 0$ for all j , but this contradict with our assumptions in 1.

(3 \Rightarrow 4) Assume that $H_i = [\langle \mathbf{v}_i, \mathbf{w} \rangle = \alpha_i]$, where $\alpha > 0$. Note that the convex hull $\Lambda_{\mathbf{W}}$ is a polytope with faces H_i . We know, $\Lambda_{\mathbf{W}}^\circ$ all lies on one side of H_i . Since $\mathbf{0} \in \Lambda_{\mathbf{W}}^\circ$,

we know for all $\mathbf{w} \in \Lambda_{\mathbf{W}}^\circ$ we have $\langle \mathbf{v}_i, \mathbf{w} \rangle < \alpha_i$, and hence $\langle \mathbf{v}_i, \tilde{\mathbf{w}}_j \rangle \leq \alpha_i$ for all j .

(4 \Rightarrow 3) is trivial.

(3 \Rightarrow 1) $\mathbf{0} \in \Lambda_{\mathbf{W}}^\circ$ implies there exist positive numbers λ_j , such that

$$\sum_{j=1}^k \lambda_j \tilde{\mathbf{w}}_j = \mathbf{0}.$$

Hence for any unit vector \mathbf{n} , we have

$$\sum_{j=1}^k \lambda_j \langle \mathbf{n}, \tilde{\mathbf{w}}_j \rangle = 0.$$

Since $\tilde{\mathbf{w}}_j$ are in general position, so $\langle \mathbf{n}, \tilde{\mathbf{w}}_j \rangle$ cannot all be 0. Hence, there must be both positive and negative terms. Now, we get the desired result. \square

Proof of Proposition 3. By Lemma 4, we know, for any $t \in T_1$, there exists a $\alpha \in [0, \frac{\pi}{2}]$ and a unit vector \mathbf{v} , such that $\langle \mathbf{v}, \tilde{\mathbf{w}}_j^{(t)} \rangle \geq \sin \alpha$ for all j and $\langle \mathbf{v}, \tilde{\mathbf{w}}_j^{(t)} \rangle = \sin \alpha$ for at least d j 's. W.l.o.g, we assume $\mathbf{v} = (1, 0, \dots, 0)$. For any $\mathbf{x} \in V_1$, we write $\tilde{\mathbf{x}} = \frac{\mathbf{x}}{|\mathbf{x}|} \in \mathcal{S}^{d-1}$. Also, we can assume for all $j \leq d$, we have $\langle \mathbf{v}, \tilde{\mathbf{w}}_j^{(t)} \rangle = \sin \alpha$.

Note that $|W|$ is bounded by R , we know there exist $\beta \in (0, \frac{\pi}{2} - \alpha)$ such that $\sum_{j=1}^k |\mathbf{w}_j| \leq R = \frac{1}{M \sin \beta}$. Now, w.l.o.g, we can assume $\tilde{\mathbf{w}}_1^{(t)} = (\cos \alpha, \sin \alpha, 0, \dots, 0)$. Take $\mathbf{n} = (-\sin \alpha, \cos \alpha, 0, \dots, 0)$, we know $\langle \mathbf{n}, \tilde{\mathbf{w}}_1^{(t)} \rangle = 0$. For all $\mathbf{x} \in V_1$ such that $\langle \mathbf{n}, \tilde{\mathbf{x}} \rangle > \cos \beta$, we have

$$\begin{aligned}
& \cos \beta \\
&< \langle \mathbf{n}, \tilde{\mathbf{x}} \rangle - \tilde{x}_1 \cos \alpha + \tilde{x}_2 \sin \alpha \\
&\leq -\tilde{x}_1 \cos \alpha + \sqrt{1 - \tilde{x}_1^2} \sin \alpha.
\end{aligned}$$

So, we have

$$\tilde{x}_1 < -\cos(\alpha + \beta).$$

Now, for all $\mathbf{x} \in V_1$ such that $\tilde{x}_1 < -\cos(\alpha + \beta)$, we have

$$\begin{aligned}
\tilde{f}(\mathbf{W}^{(t)}, \mathbf{x}) &\leq M \sum_{j=0}^d \sigma \left(\langle \mathbf{w}_j^{(t)}, \tilde{\mathbf{x}} \rangle \right) \\
&\leq M \sum_{j=0}^d |\mathbf{w}_j^{(t)}| \sin \beta < 1.
\end{aligned}$$

So, with

$$D_{\mathbf{W}^{(t)}}^1 := \{\mathbf{x} \in V_1 : \langle \mathbf{n}, \tilde{\mathbf{x}} \rangle > \cos \beta \text{ and } \langle \tilde{\mathbf{w}}_1, \tilde{\mathbf{x}} \rangle > 0\},$$

we have

$$D_{\mathbf{W}^{(t)}}^1 \subset \Omega_{\mathbf{W}^{(t)}}^1.$$

See Fig 6 for a intuition of $D_{\mathbf{W}^{(t)}}^1$. Note that \mathbf{n} and $\tilde{\mathbf{w}}_1$ perpendicular to each other, and the probability distribution p_1 is rotational invariant, so there exists a constant depending only on R , such that

$$\mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}}^1 \right] \geq \mathbb{P} \left[D_{\mathbf{W}^{(t)}}^1 \right] = p_R > 0,$$




Figure 6: 2-dim section of \mathbb{R}^d spanned by \tilde{w}_1 and n

where we have the following estimate for p_R

$$\begin{aligned}
p_R &= \int_{\mathcal{D}_{\mathbf{W}^{(t)}}^1} \langle \tilde{w}_1^{(t)}, \nabla_{w_1} l(\mathbf{W}) \rangle p_1(\mathbf{x}) d\mathbf{x} \\
&\geq \frac{p_{\min} |\mathcal{S}^{d-2}|}{|\mathcal{S}^{d-1}|} \int_{\cos \beta}^1 (1-y^2)^{\frac{d-3}{2}} dy \\
&= \frac{p_{\min} |\mathcal{S}^{d-2}|}{|\mathcal{S}^{d-1}|} \int_0^\beta (\sin \theta)^{d-2} d\theta \\
&= \Omega \left(\frac{p_{\min} (\sin \beta)^{d-1} |\mathcal{S}^{d-2}|}{(d-1) |\mathcal{S}^{d-1}|} \right) \\
&= \Omega \left(\frac{p_{\min}}{\sqrt{d} (MR)^d} \right).
\end{aligned}$$

Now, we know the gradient of w_1 on \tilde{w}_1 direction is bounded below, with same arguments in proof of Lemma 3, we have

$$\left| \langle \tilde{w}_1^{(t)}, \nabla_{w_1} l(\mathbf{W}^{(t)}) \rangle \right| \geq \frac{p_R}{2C_p}.$$

Note that

$$\left| w_1^{(t+1)} \right| - \left| w_1^{(t)} \right| \geq \eta \left| \langle \tilde{w}_1^{(t)}, \nabla_{w_1} l(\mathbf{W}^{(t)}) \rangle \right| \geq \frac{\eta p_R}{2C_p}$$

and since $\left| w_1^{(t)} \right|$ is non-decreasing, we have

$$\sum_{t=1}^{\infty} \left| w_1^{(t+1)} \right| - \left| w_1^{(t)} \right| \leq R.$$

Combining the above two equations, it follows $|T_1| \leq \frac{2C_p R}{\eta p_R}$. \square

Proof of Proposition 4. By Lemma 4, we know that for any $t \in T_2$

$$\Omega_{\mathbf{W}^{(t)}} = \bigcup_{j=1}^k \Omega_{\mathbf{W}^{(t)}}^j,$$

and thus

$$\sum_{j=1}^k \mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}}^j \right] \geq \mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}} \right].$$

On the other hand, by the same arguments in proof of Lemma 3, we know

$$\left\langle \tilde{w}_j^{(t)}, \nabla_{w_j} l(\mathbf{W}^{(t)}) \right\rangle \geq \frac{\mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}}^j \right]}{2C_p}.$$

So,

$$\left| w_j^{(t+1)} \right| - \left| w_j^{(t)} \right| \geq \eta \left\langle \tilde{w}_j^{(t)}, \nabla_{w_j} l(\mathbf{W}^{(t)}) \right\rangle \geq \frac{\eta \mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}}^j \right]}{2C_p}.$$

On the other hand, since $|\mathbf{W}^{(t)}|$ is bounded and $\mathbf{x} \in \mathcal{X}$ is bounded almost surely, we know $l(\mathbf{W}, \mathbf{x})$ is bounded by RM , then

$$l(\mathbf{W}^{(t)}) = \int_{\Omega_{\mathbf{W}^{(t)}}} l(\mathbf{W}^{(t)}, \mathbf{x}) d\mathbf{x} \leq RM \mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}} \right].$$

Combining the above 2 equations, we get

$$\left| w_j^{(t+1)} \right| - \left| w_j^{(t)} \right| \geq \frac{\eta l(\mathbf{W}^{(t)})}{2RMC_p} \frac{\mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}}^j \right]}{\mathbb{P} \left[\Omega_{\mathbf{W}^{(t)}} \right]}.$$

Note that W is bounded, we know

$$R \geq \sup_t |\mathbf{W}^t| \geq \sum_{t \in T_2} \sum_{j=1}^k \left| w_j^{(t+1)} \right| - \left| w_j^{(t)} \right| \geq \sum_{t \in T_2} \frac{\eta l(\mathbf{W}^{(t)})}{2RMC_p}.$$

Now, we get the desired result. \square