
A GLOBAL CONVERGENCE THEORY FOR DEEP RELU IMPLICIT NETWORKS VIA OVER-PARAMETERIZATION

Tianxiang Gao

Department of Computer Science
Iowa State University
gaotx@iastate.edu

Hailiang Liu

Department of Mathematics
Iowa State University
hliu@iastate.edu

Jia Liu

Dept. of Electrical and Computer Engineering
The Ohio State University
liu@ece.osu.edu

Hridesh Rajan

Department of Computer Science
Iowa State University
hridesh@iastate.edu

Hongyang Gao

Department of Computer Science
Iowa State University
hygao@iastate.edu

ABSTRACT

Implicit deep learning has received increasing attention recently, since it generalizes the recursive prediction rules of many commonly used neural network architectures. Its prediction rule is provided implicitly based on the solution of an equilibrium equation. Although many recent studies have experimentally demonstrated its superior performances, the theoretical understanding of implicit neural networks is limited. In general, the equilibrium equation may not be well-posed during the training. As a result, there is no guarantee that a vanilla (stochastic) gradient descent (SGD) training nonlinear implicit neural networks can converge. This paper fills the gap by analyzing the gradient flow of Rectified Linear Unit (ReLU) activated implicit neural networks. For an m -width implicit neural network with ReLU activation and n training samples, we show that a randomly initialized gradient descent converges to a global minimum at a linear rate for the square loss function if the implicit neural network is *over-parameterized*. It is worth noting that, unlike existing works on the convergence of (S)GD on finite-layer over-parameterized neural networks, our convergence results hold for implicit neural networks, where the number of layers is *infinite*.

1 INTRODUCTION

1) Background and Motivation: In the last decade, implicit deep learning (El Ghaoui et al., 2019) have attracted more and more attention. Its popularity is mainly because it generalizes the recursive rules of many widely used neural network architectures. A line of recent works (Bai et al., 2019; El Ghaoui et al., 2019; Bai et al., 2020) have shown that the implicit neural network architecture is a wider class that includes most current neural network architectures as special cases, such as feed-forward neural networks, convolution neural networks, residual networks, and recurrent neural networks. Moreover, implicit deep learning is also well known for its competitive performance compared to other regular deep neural networks but using significantly fewer computational resources (Dabre & Fujita, 2019; Dehghani et al., 2018; Bai et al., 2018).

Although a line of literature has been shown the superior performance of implicit neural networks experimentally, the theoretical understanding is still limited. To date, it is still unknown if a simple first-order optimization method such as (stochastic) gradient descent can converge on an implicit neural network activated by a nonlinear function. Unlike a regular deep neural network, an implicit neural network could have infinitely many layers, resulting in the possibility of divergence of the

forward propagation (El Ghaoui et al., 2019; Kawaguchi, 2021). The main challenge in establishing the convergence of implicit neural network training lies in the fact that, in general, the equilibrium equation of implicit neural networks cannot be solved in closed-form. What exacerbates the problem is the well-posedness of the forward propagation. In other words, the equilibrium equation may have zero or multiple solutions. A line of recent studies have suggested a number of strategies to handle this well-posedness challenge, but they all involved reformulation or solving subproblems in each iteration. For example, El Ghaoui et al. (2019) suggested to reformulate the training as the Fenchel divergence formulation and solve the reformulated optimization problem by projected gradient descent method. However, this requires solving a projection subproblem in each iteration and convergence was only demonstrated numerically. By using an extra softmax layer, Kawaguchi (2021) established global convergence result of gradient descent for a linear implicit neural network. Unfortunately, their result cannot be extended to nonlinear activations, which are critical to the learnability of deep neural networks.

This paper proposes a global convergence theory of gradient descent for implicit neural networks activated by nonlinear Rectified Linear Unit (ReLU) activation function by using overparameterization. Specifically, we show that random initialized gradient descent with fixed stepsize converges to a global minimum of a ReLU implicit neural network at a linear rate as long as the implicit neural network is *overparameterized*. Recently, over-parameterization has been shown to be effective in optimizing finite-depth neural networks (Zou et al., 2020; Nguyen & Mondelli, 2020; Arora et al., 2019; Oymak & Soltanolkotabi, 2020). Although the objective function in the training is non-smooth and non-convex, it can be shown that GD or SGD converge to a global minimum linearly if the width m of each layer is a polynomial of the number of training sample n and the number of layers h , *i.e.*, $m = \text{poly}(n, h)$. However, these results cannot be directly applied to implicit neural networks, since implicit neural networks have infinitely many hidden layers, *i.e.*, $h \rightarrow \infty$, and the well-posedness problem surfaces during the training process. In fact, Chen et al. (2018); Bai et al. (2019; 2020) have all observed that the time and number of iterations spent on forward propagation are gradually increased with the the training epochs. Thus, we have to ensure the unique equilibrium point always exists throughout the training given that the width m is only polynomial of n .

2) Preliminaries of Implicit Deep Learning: In this work, we consider an implicit neural network with the transition at the ℓ -th layer in the following form (El Ghaoui et al., 2019; Bai et al., 2019):

$$\mathbf{z}^\ell = \sigma \left(\frac{\gamma}{\sqrt{m}} \mathbf{A} \mathbf{z}^{\ell-1} + \phi(\mathbf{x}) \right), \quad (1)$$

where $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is a feature mapping function that transforms an input vector $\mathbf{x} \in \mathbb{R}^d$ to a desired feature vector $\phi \triangleq \phi(\mathbf{x})$, $\mathbf{z}^\ell \in \mathbb{R}^m$ is the output of the ℓ -th layer, $\mathbf{A} \in \mathbb{R}^{m \times m}$ is a trainable weight matrix, $\sigma(u) = \max\{0, u\}$ is the ReLU activation function, and $\gamma \in (0, 1)$ is a fixed scalar to scale \mathbf{A} . As will be shown later in Section 2.1, γ plays the role of ensuring the existence of the limit $\mathbf{z}^* = \lim_{\ell \rightarrow \infty} \mathbf{z}^\ell$. In general, the feature mapping function ϕ is a nonlinear function, which extracts features from the low-dimensional input vector \mathbf{x} . In this paper, we consider a simple nonlinear feature mapping function ϕ given by

$$\phi(\mathbf{x}) \triangleq \frac{1}{\sqrt{m}} \sigma(\mathbf{W} \mathbf{x}), \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{m \times d}$ is a trainable parameter matrix. As $\ell \rightarrow \infty$, an implicit neural network can be considered an *infinitely* deep neural network. Consequently, \mathbf{z}^* is not only the limit of the sequence $\{\mathbf{z}^\ell\}_{\ell=0}^\infty$ with $\mathbf{z}_0 = \mathbf{0}$, but it is also the *equilibrium point* (or *fixed point*) of the equilibrium equation:

$$\mathbf{z}^* = \sigma(\tilde{\gamma} \mathbf{A} \mathbf{z}^* + \phi), \quad (3)$$

where $\tilde{\gamma} \triangleq \gamma/\sqrt{m}$. In implicit neural networks, the prediction \hat{y} for the input vector \mathbf{x} is the combination of the fixed point \mathbf{z}^* and the feature vector ϕ , *i.e.*,

$$\hat{y} = \mathbf{u}^T \mathbf{z}^* + \mathbf{v}^T \phi, \quad (4)$$

where $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$ are trainable weight vectors. For simplicity, we use $\boldsymbol{\theta} \triangleq \text{vec}(\mathbf{A}, \mathbf{W}, \mathbf{u}, \mathbf{v})$ to group all training parameters. Given a training data set $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, we want to minimize

$$L(\boldsymbol{\theta}) = \sum_{i=1}^n \frac{1}{2} (\hat{y}_i - y_i)^2 = \frac{1}{2} \|\hat{\mathbf{y}} - \mathbf{y}\|^2, \quad (5)$$

where $\hat{\mathbf{y}}$ and \mathbf{y} are the vectors formed by stacking all the prediction and labels.

3) Main Results: Our results are based on the following observations. We first analyze the forward propagation and find that the unique equilibrium point always exists if the scaled matrix $\tilde{\gamma}\mathbf{A}$ in Eq. (3) has an operator norm less than one. Thus, the well-posedness problem is reduced to finding a sequence of scalars $\{\gamma_k\}_{k=1}^{\infty}$ such that $\tilde{\gamma}_k\mathbf{A}(k)$ is appropriately scaled. To achieve this goal, we show that the operator norm $\mathbf{A}(k)$ is uniformly upper bounded by a constant over all iterations. Consequently, a fixed scalar γ is enough to ensure the well-posedness of Eq. (3). Our second observation is from the analysis of the gradient descent method with infinitesimal step-size (gradient flow). By applying the chain rule with the gradient flow, we derive the dynamics of prediction $\hat{\mathbf{y}}(t)$ which is governed by the spectral property of a Gram matrix. In particular, if the smallest eigenvalue of the Gram matrix is lower bounded throughout the training, the gradient descent method enjoys a linear convergence rate. Along with some basic functional analysis results, it can be shown that the smallest eigenvalue of the Gram matrix at initialization is lower bounded if no two data samples are parallel. Although the Gram matrix varies in each iteration, the spectral property is preserved if the Gram matrix is close to its initialization. Thus, the convergence problem is reduced to showing the Gram matrix in latter iterations is close to its initialization. Our third observation is that we find random initialization, over-parameterization, and linear convergence jointly enforce the (operator) norms of parameters upper bounded by some constants and close to their initialization. Accordingly, we can use this property to show that the operator norm of \mathbf{A} is upper bounded and the spectral property of the Gram matrix is preserved throughout the training. Combining all these insights together, we can conclude that the random initialized gradient descent method with a constant step-size converges to a global minimum of the implicit neural network with ReLU activation.

The main contributions of this paper are summarized as follows:

- (i) By scaling the weight matrix \mathbf{A} with a fixed scalar γ , we show that the unique equilibrium point \mathbf{z}^* for each \mathbf{x} always exists during the training if the parameters are randomly initialized, even for the nonlinear ReLU activation function.
- (ii) We analyze the gradient flow of implicit neural networks. Despite the non-smooth and non-convexity of the objective function, the convergence to a global minimum at a linear rate is guaranteed if the implicit neural network is over-parameterized and the data is non-degenerate.
- (iii) Since gradient descent is discretized version of gradient flow, we can show gradient descent with fixed stepsize converges to a global minimum of implicit neural networks at a linear rate under the same assumptions made by the gradient flow analysis, as long as the stepsize is chosen small enough.

Notation: For a vector \mathbf{x} , $\|\mathbf{x}\|$ is the Euclidean norm of \mathbf{x} . For a matrix \mathbf{A} , $\|\mathbf{A}\|$ is the operator norm of \mathbf{A} . If \mathbf{A} is a square matrix, then $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ denote the smallest and largest eigenvalue of \mathbf{A} , respectively, and $\lambda_{\max}(\mathbf{A}) \leq \|\mathbf{A}\|$. We denote $[n] \triangleq \{1, 2, \dots, n\}$.

2 WELL-POSEDNESS AND GRADIENT COMPUTATION

In this section, we provide a simple condition for the equilibrium equation (3) to be well-posed in the sense that the unique equilibrium point exists. Instead of backpropagating through all the intermediate iterations of a forward pass, we derive the gradients of trainable parameters by using the implicit function theorem. In this work, we make the following assumption on parameter initialization.

Assumption 1 (Random Initialization). The entries \mathbf{A}_{ij} and \mathbf{W}_{ij} are randomly initialized by the standard Gaussian distribution $\mathcal{N}(0, 1)$, and \mathbf{u}_i and \mathbf{v}_i are randomly initialized by the symmetric Bernoulli or Rademacher distribution.

Remark 2.1. This initialization is similar to the approaches widely used in practice (Glorot & Bengio, 2010; He et al., 2015). The result obtained in this work can be easily extended to the case where the distributions for \mathbf{A}_{ij} , \mathbf{W}_{ij} , \mathbf{u}_i , and \mathbf{v}_i are replaced by *sub-Gaussian random variables*.

2.1 FORWARD PROPAGATION AND WELL-POSEDNESS

In a general implicit neural network, Eq. (3) is not necessarily well-posed, since it may admit zero or multiple solutions. In this work, we show that scaling the matrix \mathbf{A} with $\tilde{\gamma} = \gamma/\sqrt{m}$ guarantees

the existence and uniqueness of the equilibrium point \mathbf{z}^* with random initialization. This follows from a foundational result in random matrix theory as restated in the following lemma.

Lemma 2.1 (Vershynin (2018), Theorem 4.4.5). Let \mathbf{A} be an $m \times n$ random matrix whose entries A_{ij} are independent, zero-mean, and sub-Gaussian random variables. Then, for any $t > 0$, we have $\|\mathbf{A}\| \leq CK(\sqrt{m} + \sqrt{n} + t)$ with probability at least $1 - 2e^{-t^2}$. Here $C > 0$ is a fixed constant, and $K = \max_{i,j} \|A_{ij}\|_{\psi_2}$.

Under Assumption 1, Lemma 2.1 implies that, with exponentially high probability, $\|\mathbf{A}\| \leq c\sqrt{m}$ for some constant $c > 0$. By scaling \mathbf{A} by a positive scalar $\tilde{\gamma}$, we show that the transition Eq. (1) is a contraction mapping. Thus, the unique equilibrium point exists with detailed proof in Appendix A.1.

Lemma 2.2. If $\|\mathbf{A}\| \leq c\sqrt{m}$ for some $c > 0$, then for any $\gamma_0 \in (0, 1)$, the scalar $\gamma \triangleq \min\{\gamma_0, \gamma_0/c\}$ uniquely determines the existence of the equilibrium \mathbf{z}^* for every \mathbf{x} , and $\|\mathbf{z}^\ell\| \leq \frac{1}{1-\gamma_0}\|\phi\|$ for all ℓ .

Lemma 2.2 indicates that equilibria always exist if we can maintain the operator norm of the scaled matrix $(\gamma/\sqrt{m})\mathbf{A}$ less than 1 during the training. However, the operator norms of matrix \mathbf{A} are changed by the update of the gradient descent. It is hard to use a fixed scalar γ to scale the matrix \mathbf{A} over all iterations, unless the operator norm of \mathbf{A} is bounded. In Section 3, we will show $\|\mathbf{A}\| \leq 2c\sqrt{m}$ always holds throughout the training, provided that $\|\mathbf{A}(0)\| \leq c\sqrt{m}$ at initialization and the width m is sufficiently large. Thus, by using the scalar $\gamma = \min\{\gamma_0, \gamma_0/(2c)\}$ for any $\gamma_0 \in (0, 1)$, equilibria always exist and the equilibrium equation Eq. (3) is well-posed.

2.2 BACKWARD GRADIENT COMPUTING

For a regular network with finite layers, one needs to store all intermediate parameters and apply backpropagation to compute the gradients. In contrast, for an implicit neural network, we can derive the formula of the gradients by using the implicit function theorem. Here, the equilibrium point \mathbf{z}^* is a root of the function f given by $f(\mathbf{z}, \mathbf{A}, \mathbf{W}) \triangleq \mathbf{z} - \sigma(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi)$. The essential challenge is to ensure the partial derivative $\partial f/\partial \mathbf{z}$ at \mathbf{z}^* is always invertible throughout the training. The following lemma shows that the partial derivative $\partial f/\partial \mathbf{z}$ at \mathbf{z}^* always exists with the scalar γ , and the gradient derived in the lemma always exists. The gradient derivations for each trainable parameters are provided in the following lemma and the proof is provided in Appendix A.2.

Lemma 2.3 (Gradients of an Implicit Neural Network). Define the scalar $\gamma \triangleq \min\{\gamma_0, \gamma_0/c\}$. If $\|\mathbf{A}\| \leq c\sqrt{m}$ for some constant $c > 0$, then for any $\gamma_0 \in (0, 1)$, the following results hold:

- (i) The partial derivatives of f with respect to \mathbf{z} , \mathbf{A} , and \mathbf{W} are

$$\frac{\partial f}{\partial \mathbf{z}} = [\mathbf{I}_m - \tilde{\gamma} \mathbf{diag}(\sigma'(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi))\mathbf{A}]^T, \quad (6)$$

$$\frac{\partial f}{\partial \mathbf{A}} = -\tilde{\gamma} [\mathbf{z}^T \otimes \mathbf{diag}(\sigma'(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi))]^T, \quad (7)$$

$$\frac{\partial f}{\partial \mathbf{W}} = -\frac{1}{\sqrt{m}} [\mathbf{x}^T \otimes \mathbf{diag}(\sigma'(\mathbf{W}\mathbf{x}))]^T \mathbf{diag}(\sigma'(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi))^T, \quad (8)$$

where $\tilde{\gamma} \triangleq \gamma/\sqrt{m}$.

- (ii) For any vector \mathbf{v} , the following inequality holds

$$\lambda_{\min}\{\mathbf{I}_m - \tilde{\gamma} \mathbf{diag}(\sigma'(\mathbf{v}))\mathbf{A}\} > 1 - \gamma_0 > 0, \quad (9)$$

which further implies $\partial f/\partial \mathbf{z}$ is invertible at the equilibrium point \mathbf{z}^* .

- (iii) The gradient of the objective function L with respect to \mathbf{A} and \mathbf{W} are given by

$$\nabla_{\mathbf{u}} L = \sum_{i=1}^n (\hat{y}_i - y_i) \mathbf{z}_i, \quad \nabla_{\mathbf{v}} L = \sum_{i=1}^n (\hat{y}_i - y_i) \phi_i, \quad (10)$$

$$\nabla_{\mathbf{A}} L = \sum_{i=1}^n \tilde{\gamma} (\hat{y}_i - y_i) \mathbf{D}_i^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-T} \mathbf{u} \mathbf{z}_i^T, \quad (11)$$

$$\nabla_{\mathbf{W}} L = \sum_{i=1}^n \frac{1}{\sqrt{m}} (\hat{y}_i - y_i) \mathbf{E}_i^T \left\{ \mathbf{D}_i^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-T} \mathbf{u} + \mathbf{v} \right\} \mathbf{x}_i^T, \quad (12)$$

where \mathbf{z}_i is the equilibrium point for the training data (\mathbf{x}_i, y_i) , $\mathbf{D}_i \triangleq \mathbf{diag}(\sigma'(\tilde{\gamma}\mathbf{A}\mathbf{z}_i + \phi_i))$, and $\mathbf{E}_i \triangleq \mathbf{diag}(\sigma'(\mathbf{W}\mathbf{x}_i))$.

3 GLOBAL CONVERGENCE OF THE GRADIENT DESCENT METHOD

In this section, we establish the global convergence results of the gradient descent method in implicit neural networks. In Section 3.1, we first study the dynamics of the prediction induced by the gradient flow, that is, the gradient descent with infinitesimal step-size. We show that the dynamics of the prediction is controlled by a Gram matrix whose smallest eigenvalue is strictly positive over iterations with high probability. Based on the findings in gradient flow, we will show that the random initialized gradient descent method with a constant step-size converges to a global minimum at a linear rate in Section 3.2.

3.1 CONTINUOUS TIME ANALYSIS

The gradient flow is equivalent to the gradient descent method with an infinitesimal step-size. Thus, the analysis of gradient flow can serve as a stepping stone towards understanding discrete-time gradient-based algorithms. Following previous works on gradient flow of different machine learning models (Saxe et al., 2013; Du et al., 2019; Arora et al., 2018; Kawaguchi, 2021) the gradient flow of implicit neural networks is given by the following ordinary differential equations:

$$\frac{d\text{vec}(\mathbf{A})}{dt} = -\frac{\partial L(t)}{\partial \mathbf{A}}, \quad \frac{d\text{vec}(\mathbf{W})}{dt} = -\frac{\partial L(t)}{\partial \mathbf{W}}, \quad \frac{d\mathbf{u}}{dt} = -\frac{\partial L(t)}{\partial \mathbf{u}}, \quad \frac{d\mathbf{v}}{dt} = -\frac{\partial L(t)}{\partial \mathbf{v}}, \quad (13)$$

where $L(t)$ represents the value of the objective function $L(\boldsymbol{\theta})$ at time t .

Our results relies on the analysis of the dynamics of prediction $\hat{\mathbf{y}}(t)$. In particular, Lemma 3.1 shows that the dynamics of prediction $\hat{\mathbf{y}}(t)$ is governed by the spectral property of a Gram matrix $\mathbf{H}(t)$.

Lemma 3.1 (Dynamics of Prediction $\hat{\mathbf{y}}(t)$). Suppose $\|\mathbf{A}(s)\| \leq c\sqrt{m}$ for all $0 \leq s \leq t$. Let $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\Phi(s) \in \mathbb{R}^{n \times m}$, and $\mathbf{Z}(s) \in \mathbb{R}^{n \times m}$ be the matrices whose rows are the training data \mathbf{x}_i , feature vectors ϕ_i , and equilibrium points \mathbf{z}_i at time s , respectively. With scalar $\gamma \triangleq \min\{\gamma_0, \gamma_0/c\}$ for any $\gamma_0 \in (0, 1)$, the dynamics of prediction $\hat{\mathbf{y}}(t)$ is given by

$$\begin{aligned} \frac{d\hat{\mathbf{y}}}{dt} &= -[(\gamma^2 \mathbf{M}(t) + \mathbf{I}_n) \circ \mathbf{Z}(t)\mathbf{Z}(t)^T + \mathbf{Q}(t) \circ \mathbf{X}\mathbf{X}^T + \Phi(t)\Phi(t)^T](\hat{\mathbf{y}} - \mathbf{y}), \\ &\triangleq -\mathbf{H}(t)(\hat{\mathbf{y}} - \mathbf{y}), \end{aligned} \quad (14)$$

where \circ is the Hadamard product (i.e., element-wise product), and matrices $\mathbf{M}(t) \in \mathbb{R}^{n \times n}$ and $\mathbf{Q}(t) \in \mathbb{R}^{n \times n}$ are defined as follows:

$$\mathbf{M}(t)_{ij} \triangleq \frac{1}{m} \mathbf{u}^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-1} \mathbf{D}_i \mathbf{D}_j^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u}, \quad (15)$$

$$\mathbf{Q}(t)_{ij} \triangleq \frac{1}{m} (\mathbf{D}_i^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-1} \mathbf{u} + \mathbf{v})^T \mathbf{E}_i \mathbf{E}_j^T (\mathbf{D}_j^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}). \quad (16)$$

Note that the matrix $\mathbf{H}(t)$ is clearly positive semidefinite since it is the sum of three positive semidefinite matrices. If there exists a constant $\lambda > 0$ such that $\lambda_{\min}\{\mathbf{H}(t)\} \geq \lambda > 0$ for all t , i.e., $\mathbf{H}(t)$ is positive definite, then the dynamics of the loss function $L(t)$ satisfies the following inequality

$$L(t) \leq \exp\{-\lambda t\}L(0),$$

which immediately indicates that the objective value $L(t)$ is consistently decreasing to zero at a geometric rate. With random initialization, we will show that $\mathbf{H}(t)$ is positive definite as long as the number of parameters m is sufficiently large and no two data points are parallel to each other. In particular, by using the nonlinearity of the feature map ϕ , we will show that the smallest eigenvalue of the Gram matrix $\mathbf{G}(t) \triangleq \Phi(t)\Phi(t)^T$ is strictly positive over all time t with high probability. As a result, the smallest eigenvalue of $\mathbf{H}(t)$ is always strictly positive.

Clearly, $\mathbf{G}(t)$ is a time-varying matrix. We first analyze its spectral property at its initialization. When $t = 0$, it follows from the definition of the feature vector ϕ in Eq. (2) that

$$\mathbf{G}(0) = \Phi(0)\Phi(0)^T = \frac{1}{m} \sigma(\mathbf{X}\mathbf{W}(0)^T) \sigma(\mathbf{X}\mathbf{W}(0)^T)^T = \frac{1}{m} \sum_{r=1}^m \sigma(\mathbf{X}\mathbf{w}_r(0)) \sigma(\mathbf{X}\mathbf{w}_r(0))^T.$$

By Assumption 1, each vector $\mathbf{w}_r(0)$ follows the standard multivariate normal distribution, *i.e.*, $\mathbf{w}_r(0) \sim N(\mathbf{0}, \mathbf{I}_d)$. By letting $m \rightarrow \infty$, we obtain the covariance matrix $\mathbf{G}^\infty \in \mathbb{R}^{n \times n}$ as follows:

$$\mathbf{G}^\infty \triangleq \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[\sigma(\mathbf{X}\mathbf{w})\sigma(\mathbf{X}\mathbf{w})^T]. \quad (17)$$

Here \mathbf{G}^∞ is a Gram matrix induced by the ReLU activation function and the random initialization. The following lemma shows that the smallest eigenvalue of \mathbf{G}^∞ is strictly positive as long as no two data points are parallel. Moreover, later in Lemmas 3.3 and 3.4, we conclude that the spectral property of \mathbf{G}^∞ is preserved in the Gram matrices $\mathbf{G}(0)$ and $\mathbf{G}(t)$ during the training, as long as the number of parameter m is sufficiently large.

Lemma 3.2. Assume $\|\mathbf{x}_i\| = 1$ for all $i \in [n]$. If $\mathbf{x}_i \not\parallel \mathbf{x}_j$ for all $i \neq j$, then $\lambda_0 \triangleq \lambda_{\min}\{\mathbf{G}^\infty\} > 0$.

Proof. The proof follows from the *Hermite Expansions* of the matrix \mathbf{G}^∞ and the complete proof is provided in Appendix A.4. \square

Assumption 2 (Training Data). Without loss of generality, we can assume that each \mathbf{x}_i is normalized to have a unit norm, *i.e.*, $\|\mathbf{x}_i\| = 1$, for all $i \in [n]$. Moreover, we assume $\mathbf{x}_i \not\parallel \mathbf{x}_j$ for all $i \neq j$.

In most real-world datasets, it is extremely rare that two data samples are parallel. If this happens, by adding some random noise perturbation to the data samples, Assumption 2 can still be satisfied. Next, we show that at the initialization, the spectral property of \mathbf{G}^∞ is preserved in $\mathbf{G}(0)$ if m is sufficiently large. Specifically, the following lemma shows that if $m = \tilde{\Omega}(n^2)$, then $\mathbf{G}(0)$ has a strictly positive smallest eigenvalue with high probability. The proof follows from the standard concentration bound for Gaussian random variables, and we relegate the proof to Appendix A.5.

Lemma 3.3. Let $\lambda_0 = \lambda_{\min}(\mathbf{G}^\infty) > 0$. If $m = \Omega\left(\frac{n^2}{\lambda_0^2} \log\left(\frac{n}{\delta}\right)\right)$, then with probability of at least $1 - \delta$, it holds that $\|\mathbf{G}(0) - \mathbf{G}^\infty\|_2 \leq \frac{\lambda_0}{4}$, and hence $\lambda_{\min}(\mathbf{G}(0)) \geq \frac{3}{4}\lambda_0$.

During training, $\mathbf{G}(t)$ is time-varying, but it can be shown to be close to $\mathbf{G}(0)$ and preserve the spectral property of \mathbf{G}^∞ , if the matrices $\mathbf{W}(t)$ has a bounded operator norm and it is not far away from $\mathbf{W}(0)$. This result is formally stated below and its proof is provided in Appendix A.6.

Lemma 3.4. Suppose $\|\mathbf{W}(0)\| \leq c\sqrt{m}$, and $\lambda_{\min}\{\mathbf{G}(0)\} \geq \frac{3}{4}\lambda_0$. For any matrix $\mathbf{W} \in \mathbb{R}^{m \times d}$ that satisfies $\|\mathbf{W}\| \leq 2c\sqrt{m}$ and $\|\mathbf{W} - \mathbf{W}(0)\| \leq \frac{\sqrt{m}\lambda_0}{16c\|\mathbf{X}\|^2} \triangleq R$, the matrix defined by $\mathbf{G} \triangleq \frac{1}{m}\sigma(\mathbf{X}\mathbf{W}^T)\sigma(\mathbf{X}\mathbf{W}^T)^T$ satisfies $\|\mathbf{G} - \mathbf{G}(0)\| \leq \frac{\lambda_0}{4}$ and $\lambda_{\min}(\mathbf{G}) \geq \frac{\lambda_0}{2}$.

The next lemma shows three facts: (1) The smallest eigenvalue of $\mathbf{G}(t)$ is strictly positive for all $t \geq 0$; (2) The objective value $L(t)$ converges to zero at a linear convergence rate; (3) The (operator) norms of all trainable parameters are upper bounded by some constants, which further implies that unique equilibrium points in matrices $\mathbf{Z}(t)$ always exist. We prove this lemma by induction, which is provided in Appendix A.7.

Lemma 3.5. Suppose that $\|\mathbf{u}(0)\| = \sqrt{m}$, $\|\mathbf{v}(0)\| = \sqrt{m}$, $\|\mathbf{W}(0)\| \leq c\sqrt{m}$, $\|\mathbf{A}(0)\| \leq c\sqrt{m}$, and $\lambda_{\min}\{\mathbf{G}(0)\} \geq \frac{3}{4}\lambda_0 > 0$. If $m = \Omega\left(\frac{c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2\right)$ and $0 < \gamma \leq \min\{\frac{1}{2}, \frac{1}{4c}\}$, then for any $t \geq 0$, the following results hold:

- (i) $\lambda_{\min}(\mathbf{G}(t)) \geq \frac{\lambda_0}{2}$,
- (ii) $\|\mathbf{u}(t)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iii) $\|\mathbf{v}(t)\| \leq \frac{8c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iv) $\|\mathbf{W}(t)\| \leq 2c\sqrt{m}$,
- (v) $\|\mathbf{A}(t)\| \leq 2c\sqrt{m}$,
- (vi) $\|\hat{\mathbf{y}}(t) - \mathbf{y}\|^2 \leq \exp\{-\lambda_0 t\} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2$.

By using simple union bounds in Lemma 2.1 and 3.5, we immediately obtain the global convergence result for the gradient flow as follows.

Theorem 3.1 (Convergence Rate of Gradient Flow). Suppose that Assumptions 1 and 2 hold. If we set the number of parameter $m = \Omega\left(\frac{n^2}{\lambda_0^2} \log\left(\frac{n}{\delta}\right)\right)$ and choose $0 < \gamma \leq \min\{\frac{1}{2}, \frac{1}{4c}\}$, then with probability at least $1 - \delta$ over the initialization, we have

$$\|\hat{\mathbf{y}}(t) - \mathbf{y}\|^2 \leq \exp\{-\lambda_0 t\} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2, \quad \forall t \geq 0.$$

This theorem establishes the global convergence of the gradient flow. Despite the nonconvexity of the objective function $L(\boldsymbol{\theta})$, Theorem 3.1 shows that if m is sufficiently large, then the objective value is consistently decreasing to zero at a geometric rate. In particular, Theorem 3.1 requires $m = \tilde{\Omega}(n^2)$, which is similar or even better than recent results for the neural network with *finite* layers (Nguyen & Mondelli, 2020; Oymak & Soltanolkotabi, 2020; Allen-Zhu et al., 2019; Zou & Gu, 2019; Du et al., 2019). In particular, previous results showed that m is a polynomial of the number of training sample n and the number of layers h , i.e., $m = \Omega(n^\alpha h^\beta)$ with $\alpha \geq 2$ and $\beta \geq 12$ (Nguyen & Mondelli, 2020, Table 1). These results do not apply in our case since we have *infinitely* many layers. By taking advantage of the nonlinear feature mapping function, we establish the global convergence for the gradient flow with m independent of depth h .

3.2 DISCRETE TIME ANALYSIS

In this section, we show that the randomly initialized gradient descent method with a fixed step-size converges to a global minimum at a linear rate. With similar argument used in the analysis of gradient flow, we can show the (operator) norms of the training parameters are upper bounded by some constants. It is worth noting that, unlike the analysis of gradient flow, we do not have an explicit formula for the dynamics of prediction $\hat{\mathbf{y}}(t)$ in the discrete time analysis. Instead, we have to show the difference between the equilibrium points $\mathbf{Z}(k)$ in two consecutive iterations are bounded. Based on this, we can further bound the changes in the predictions $\hat{\mathbf{y}}(k)$ between two consecutive iterations. Another challenge is to show the objective value consistently decreases over iterations. A general strategy is to show the objective function is (semi-)smooth with respect to parameter $\boldsymbol{\theta}$, and apply the descent lemma to the objective function (Nguyen & Mondelli, 2020; Allen-Zhu et al., 2019; Zou & Gu, 2019). In this section, we take advantage of the nonlinear feature mapping function in Eq. (4). Consequently, we are able to obtain a Polyak-Łojasiewicz-like condition (Karimi et al., 2016; Nguyen, 2021), which allows us to provide a much simpler proof.

The following lemma establishes the global convergence for the gradient descent method with a fixed step-size when the operator norms of $\mathbf{A}(0)$ and $\mathbf{W}(0)$ are bounded and $\lambda_{\min}(\mathbf{G}(0)) > 0$. The proof is proved in Appendix A.8.

Lemma 3.6 (Gradient Descent Convergence Rate). Suppose $\|\mathbf{u}(0)\| = \sqrt{m}$, $\|\mathbf{v}(0)\| = \sqrt{m}$, $\|\mathbf{W}(0)\| \leq c\sqrt{m}$, $\|\mathbf{A}(0)\| \leq c\sqrt{m}$, and $\lambda_{\min}\{\mathbf{G}(0)\} \geq \frac{3}{4}\lambda_0 > 0$. If $0 < \gamma \leq \min\{\frac{1}{2}, \frac{1}{4c}\}$, $m = \Omega\left(\frac{c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2\right)$, and stepsize $\alpha = \mathcal{O}(\lambda_0/n^2)$, then for any $k \geq 0$, we have

- (i) $\lambda_{\min}(\mathbf{G}(k)) \geq \frac{\lambda_0}{2}$,
- (ii) $\|\mathbf{u}(k)\| \leq \frac{32c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iii) $\|\mathbf{v}(k)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iv) $\|\mathbf{W}(k)\| \leq 2c\sqrt{m}$,
- (v) $\|\mathbf{A}(k)\| \leq 2c\sqrt{m}$,
- (vi) $\|\hat{\mathbf{y}}(k) - \mathbf{y}\|^2 \leq (1 - \alpha\lambda_0/2)^k \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2$.

By using simple union bounds to combine Lemma 2.1 and 3.6, we obtain the global convergence result for the gradient descent.

Theorem 3.2 (Convergence Rate of Gradient Descent). Suppose that Assumption 1 and 2 hold. If we set $m = \Omega\left(\frac{n^2}{\lambda_0} \log\left(\frac{n}{\delta}\right)\right)$, $0 < \gamma \leq \min\{\frac{1}{2}, \frac{1}{4c}\}$, and choose step-size $\alpha = \mathcal{O}(\lambda_0/n^2)$, then with probability at least $1 - \delta$ over the initialization, we have

$$\|\hat{\mathbf{y}}(k) - \mathbf{y}\|^2 \leq (1 - \alpha\lambda_0/2)^k \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2, \quad \forall k \geq 0.$$

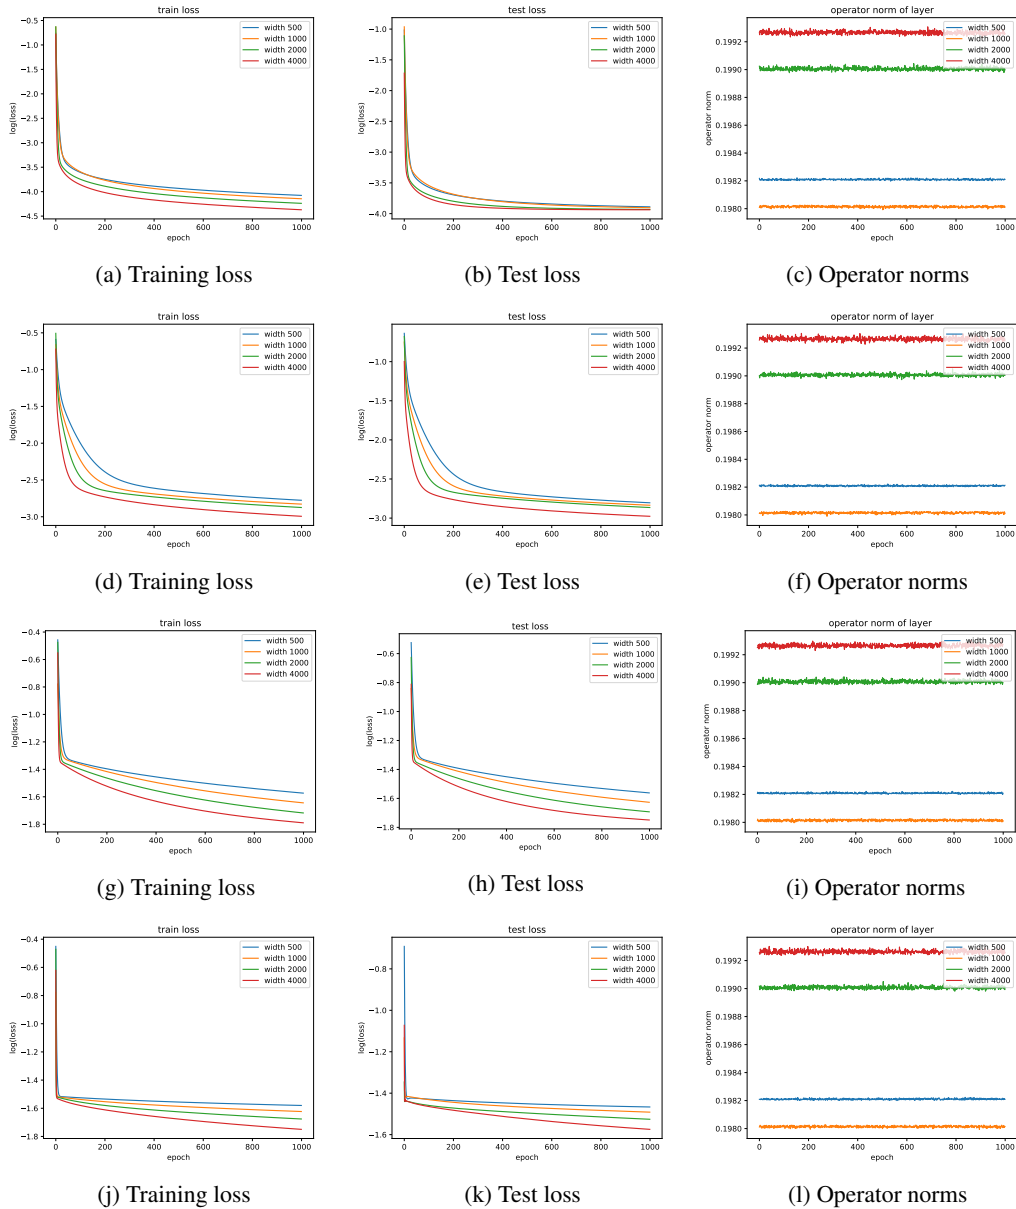


Figure 1: Results on MNIST, FashionMNIST, CIFAR10, and SVHN. We evaluate the impact of the width m on the training loss, test loss, and operator norm of the scaled matrix $(\gamma/\sqrt{m})\mathbf{A}(k)$ on four real datasets.

4 EXPERIMENTAL RESULTS

In this section, we use real-world datasets MNST, FashionMNIST, CIFAR10, and SVHN to evaluate our theoretical findings. We initialize the entries of parameters \mathbf{A} , \mathbf{W} , \mathbf{u} , and \mathbf{v} by standard Gaussian or symmetric Bernoulli distribution independently as suggested in Assumption 1. For each dataset, we only use classes 0 and 1, and 500 samples are randomly drawn from each class to generate the training dataset with $n = 1000$ samples. All data samples are converted to gray scale and resized into a 28×28 pixel image. We also normalize each data to have unit norm. If two parallel samples are observed, we add a random Gaussian noise perturbation to one of them. Thus, Assumption 2 is also satisfied. We run 1000 epochs of gradient descent with a fixed step-size. We test three metrics with different width m . We first test how the extent of over-parameterization

affects the convergence rates. Then, we test the relation between the extent of over-parameterization and the operator norms between matrix $\mathbf{A}(k)$ and its initialization. Note that the “operator norm” in the plots denotes the operator norm of the scaled matrix $(\gamma/\sqrt{m})\|\mathbf{A}(k)\|$. Third, we test the extent of over-parameterization and the performance of the trained neural network on the unseen test data. Similar to the training dataset, We randomly select 500 samples from each class as the test dataset.

From Figure 1, the figures in the first column show that as m becomes larger, we have better convergence rates. The figures in the second column show that as m becomes larger, the neural networks achieve lower test loss. The figures in the third column show that the operator norms are slightly larger for larger m but overall the operator norms are approximately equal to its initialization. The bell curve from the classical bias-variance trade-off does not appear. This opens the door to a new research direction in implicit neural networks in the analyses of generalization error and bias-variance trade-off.

5 RELATED WORKS

Implicit models has been explored explored by the deep learning community for decades. For example, Pineda (1987) and ALMEIDA (1987) studied implicit differentiation techniques for training recurrent dynamics, also known as *recurrent back-propagation* (Liao et al., 2018). Recently, there has been renewed interested in the implicit models in the deep learning community (El Ghaoui et al., 2019; Gould et al., 2019). For example, Bai et al. (2019) introduces an implicit neural network called *deep equilibrium model* for the for the task of sequence modeling. By using implicit ODE solvers, Chen et al. (2018) proposed *neural ordinary differential equation* as an implicit residual network with continuous-depth. Other instantiations of implicit modeling include optimization layers (Djulonga & Krause, 2017; Amos & Kolter, 2017), differentiable physics engines (de Avila Belbute-Peres et al., 2018; Qiao et al., 2020), logical structure learning (Wang et al., 2019), and continuous generative models (Grathwohl et al., 2019).

The theoretical study of training finite-layer neural networks via over-parameterization has been an active research area. Jacot et al. (2018) showed the trajectory of the gradient descent method can be characterized by a kernel called *neural tangent kernel* for smooth activation and infinitely width neural networks. For a finite-width neural network with smooth or ReLU activation, Arora et al. (2019); Du et al. (2019); Li & Liang (2018) showed that the dynamics of the neural network is governed by a Gram matrix. Consequently, Zou et al. (2020); Du et al. (2019); Allen-Zhu et al. (2019); Nguyen & Mondelli (2020); Arora et al. (2019); Zou & Gu (2019); Oymak & Soltanolkotabi (2020) showed that (stochastic) gradient descent can attain global convergence for training a finite-layer neural network when the width m is a polynomial of the sample size n and the depth h . However, their results cannot be applied directly to implicit neural networks since implicit neural networks have infinite layers and the equilibrium equation may not be well-posed. Our work establishes the well-posedness of the equilibrium equation even if the width m is only square of the sample size n .

6 CONCLUSION AND FUTURE WORK

In this paper, we provided a convergence theory for implicit neural networks with ReLU activation in the over-parameterization regime. We showed that the random initialized gradient descent method with fixed step-size converges to a global minimum of the loss function at a linear rate if the width $m = \tilde{\Omega}(n^2)$. In particular, by using a fixed scalar $\gamma \in (0, 1)$ to scale the random initialized weight matrix \mathbf{A} , we proved that the equilibrium equation is always well-posed throughout the training. By analyzing the gradient flow, we observe that the dynamics of the prediction vector is controlled by a Gram matrix whose smallest eigenvalue is lower bounded by a strictly positive constant as long as $m = \tilde{\Omega}(n^2)$. We envision several potential future directions based on the observations made in the experiments. First, we believe that our analysis can be generalized to implicit neural networks with other scaling techniques and initialization. Here, we use a scalar γ/\sqrt{m} to ensure the existence of the equilibrium point z^* with random initialization. With an appropriate normalization, the global convergence for identity initialization can be obtained. Second, we believe that the width m not only improves the convergence rates but also the generalization performance. In particular, our experimental results showed that with a larger m value, the test loss is reduced while the classical bell curve of bias-variance trade-off is not observed.

ACKNOWLEDGMENTS

This work has been supported in part by National Science Foundation grants III-2104797, DMS1812666, CAREER CNS-2110259, CNS-2112471, CNS-2102233, CCF-2110252, CNS-21-20448, CCF-19-34884 and a Google Faculty Research Award.

REFERENCES

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252. PMLR, 2019.
- LB ALMEIDA. A learning rule for asynchronous perceptrons with feedback in a combinatorial environment. In *Proceedings, 1st First International Conference on Neural Networks*, volume 2, pp. 609–618. IEEE, 1987.
- Brandon Amos and J Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pp. 136–145. PMLR, 2017.
- Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*, 2018.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *arXiv preprint arXiv:1904.11955*, 2019.
- Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Trellis networks for sequence modeling. *arXiv preprint arXiv:1810.06682*, 2018.
- Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. *arXiv preprint arXiv:1909.01377*, 2019.
- Shaojie Bai, Vladlen Koltun, and J Zico Kolter. Multiscale deep equilibrium models. *arXiv preprint arXiv:2006.08656*, 2020.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *arXiv preprint arXiv:1806.07366*, 2018.
- Raj Dabre and Atsushi Fujita. Recurrent stacking of layers for compact neural machine translation models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 6292–6299, 2019.
- Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. End-to-end differentiable physics for learning and control. *Advances in neural information processing systems*, 31:7178–7189, 2018.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.
- Josip Djolonga and Andreas Krause. Differentiable learning of submodular models. *Advances in Neural Information Processing Systems*, 30:1013–1023, 2017.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pp. 1675–1685. PMLR, 2019.
- Laurent El Ghaoui, Fangda Gu, Bertrand Travacca, Armin Askari, and Alicia Y Tsai. Implicit deep learning. *arXiv preprint arXiv:1908.06315*, 2, 2019.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.
- Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks: A new hope. *arXiv preprint arXiv:1909.04866*, 2019.

-
- Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Fjord: Free-form continuous dynamics for scalable reversible generative models. *International Conference on Learning Representations*, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.
- Kenji Kawaguchi. On the theory of implicit deep learning: Global convergence with implicit layers. *arXiv preprint arXiv:2102.07346*, 2021.
- Erwin Kreyszig. *Introductory functional analysis with applications*, volume 1. wiley New York, 1978.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. *arXiv preprint arXiv:1808.01204*, 2018.
- Renjie Liao, Yuwen Xiong, Ethan Fetaya, Lisa Zhang, KiJung Yoon, Xaq Pitkow, Raquel Urtasun, and Richard Zemel. Reviving and improving recurrent back-propagation. In *International Conference on Machine Learning*, pp. 3082–3091. PMLR, 2018.
- Barbara MacCluer. *Elementary functional analysis*, volume 253. Springer Science & Business Media, 2008.
- Quynh Nguyen. On the proof of global convergence of gradient descent for deep relu networks with linear widths. *arXiv preprint arXiv:2101.09612*, 2021.
- Quynh Nguyen and Marco Mondelli. Global convergence of deep networks with one wide layer followed by pyramidal topology. *arXiv preprint arXiv:2002.07867*, 2020.
- Samet Oymak and Mahdi Soltanolkotabi. Toward moderate overparameterization: Global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 1(1):84–105, 2020.
- Fernando Pineda. Generalization of back propagation to recurrent and higher order neural networks. In *Neural information processing systems*, pp. 602–611, 1987.
- Yi-Ling Qiao, Junbang Liang, Vladlen Koltun, and Ming C Lin. Scalable differentiable physics for learning and control. *arXiv preprint arXiv:2007.02168*, 2020.
- Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Number 47 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018. ISBN 978-1-108-41519-4.
- Po-Wei Wang, Priya Donti, Bryan Wilder, and Zico Kolter. Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *International Conference on Machine Learning*, pp. 6545–6554. PMLR, 2019.
- Difan Zou and Quanquan Gu. An improved analysis of training over-parameterized deep neural networks. *arXiv preprint arXiv:1906.04688*, 2019.
- Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes over-parameterized deep relu networks. *Machine Learning*, 109(3):467–492, 2020.

A APPENDIX

A.1 PROOF OF LEMMA 2.2

Proof. Let $\gamma_0 \in (0, 1)$. Set $\gamma \triangleq \min\{\gamma_0, \gamma_0/c\}$. Denote $\tilde{\gamma} = \gamma/\sqrt{m}$. Then

$$\begin{aligned} \|\mathbf{z}^{\ell+1} - \mathbf{z}^\ell\| &= \|\sigma(\tilde{\gamma}\mathbf{A}\mathbf{z}^\ell + \phi) - \sigma(\tilde{\gamma}\mathbf{A}\mathbf{z}^{\ell-1} + \phi)\| \\ &\leq \tilde{\gamma} \|\mathbf{A}\mathbf{z}^\ell - \mathbf{A}\mathbf{z}^{\ell-1}\|, \quad \sigma \text{ is 1-Lipschitz continuous} \\ &= \tilde{\gamma} \|\mathbf{A}(\mathbf{z}^\ell - \mathbf{z}^{\ell-1})\| \\ &\leq \tilde{\gamma} \|\mathbf{A}\| \|\mathbf{z}^\ell - \mathbf{z}^{\ell-1}\|, \\ &\leq \tilde{\gamma} c \sqrt{m} \|\mathbf{z}^\ell - \mathbf{z}^{\ell-1}\|, \quad \|\mathbf{A}\| \leq c\sqrt{m} \\ &= \gamma_0 \|\mathbf{z}^\ell - \mathbf{z}^{\ell-1}\|. \end{aligned}$$

Applying the above argument ℓ times, we obtain

$$\|\mathbf{z}^{\ell+1} - \mathbf{z}^\ell\| \leq \gamma_0^\ell \|\mathbf{z}^1 - \mathbf{z}^0\| = \gamma_0^\ell \|\mathbf{z}^1\| = \gamma_0^\ell \|\sigma(\phi)\| \leq \gamma_0^\ell \|\phi\|,$$

where we use the fact $\mathbf{z}^0 = \mathbf{0}$. For any positive integers p, q with $p \leq q$, we have

$$\begin{aligned} \|\mathbf{z}^p - \mathbf{z}^q\| &\leq \|\mathbf{z}^p - \mathbf{z}^{p+1}\| + \dots + \|\mathbf{z}^{q-1} - \mathbf{z}^q\| \\ &\leq \gamma_0^p \|\phi\| + \dots + \gamma_0^q \|\phi\| \\ &\leq \gamma_0^p \|\phi\| (1 + \gamma_0 + \gamma_0^2 + \dots) \\ &= \frac{\gamma_0^p}{1 - \gamma_0} \|\phi\|. \end{aligned}$$

Since $\gamma_0 \in (0, 1)$, we have $\|\mathbf{z}^p - \mathbf{z}^q\| \rightarrow 0$ as $p \rightarrow \infty$. Hence, $\{\mathbf{z}^\ell\}_{\ell=1}^\infty$ is a Cauchy sequence. Since \mathbb{R}^m is complete, the equilibrium point \mathbf{z}^* is the limit of the sequence $\{\mathbf{z}^\ell\}_{\ell=1}^\infty$, so that \mathbf{z} exists and is unique. Moreover, let $q \rightarrow \infty$, then we obtain $\|\mathbf{z}^p - \mathbf{z}^*\| \leq \frac{\gamma_0^p}{1 - \gamma_0} \|\phi\|$, so that the fixed-point iteration converges to \mathbf{z} linearly.

Let $p = 0$ and $q = \ell$, then we obtain $\|\mathbf{z}^\ell\| \leq \frac{1}{1 - \gamma_0} \|\phi\|$.

□

A.2 PROOF OF LEMMA 2.3

Proof. (i) To simplify the notations, we denote $\mathbf{D} \triangleq \text{diag}(\sigma'(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi))$, and $\mathbf{E} \triangleq \text{diag}(\sigma'(\mathbf{W}\mathbf{x}))$. The differential of f is given by

$$\begin{aligned} df &= d(\mathbf{z} - \tilde{\gamma}\sigma(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi)) \\ &= d\mathbf{z} - \mathbf{D}d(\tilde{\gamma}\mathbf{A}\mathbf{z} + \phi) \\ &= [\mathbf{I}_m - \tilde{\gamma}\mathbf{D}\mathbf{A}]d\mathbf{z} - \tilde{\gamma}\mathbf{D}(d\mathbf{A})\mathbf{z} - \mathbf{D}d\phi. \end{aligned}$$

Taking vectorization on both sides yields

$$\begin{aligned} \text{vec}(df) &= [\mathbf{I}_m - \tilde{\gamma}\mathbf{D}\mathbf{A}] \text{vec}(d\mathbf{z}) - \text{vec}(\tilde{\gamma}\mathbf{D}d\mathbf{A}\mathbf{z}) - \mathbf{D}\text{vec}(d\phi) \\ &= [\mathbf{I}_m - \tilde{\gamma}\mathbf{D}\mathbf{A}] \text{vec}(d\mathbf{z}) - \tilde{\gamma}[\mathbf{z}^T \otimes \mathbf{D}] \text{vec}(d\mathbf{A}) - \mathbf{D}\text{vec}(d\phi). \end{aligned}$$

Therefore, the partial derivative of f with respect to \mathbf{z} , \mathbf{A} , and ϕ are given by

$$\begin{aligned} \frac{\partial f}{\partial \mathbf{z}} &= [\mathbf{I}_m - \tilde{\gamma}\mathbf{D}\mathbf{A}]^T, \\ \frac{\partial f}{\partial \mathbf{A}} &= -\tilde{\gamma}[\mathbf{z}^T \otimes \mathbf{D}]^T, \\ \frac{\partial f}{\partial \phi} &= -\mathbf{D}^T. \end{aligned}$$

It follows from the definition of the feature vector ϕ in Eq. (2) that

$$d\phi = \frac{1}{\sqrt{m}} d\sigma(\mathbf{W}\mathbf{x}) = \frac{1}{\sqrt{m}} \mathbf{E}(d\mathbf{W})\mathbf{x} = \frac{1}{\sqrt{m}} [\mathbf{x}^T \otimes \mathbf{E}] \text{vec}(d\mathbf{W}).$$

Thus, the partial derivative of ϕ with respect to \mathbf{W} is given by

$$\frac{\partial \phi}{\partial \mathbf{W}} = \frac{1}{\sqrt{m}} [\mathbf{x}^T \otimes \mathbf{E}]^T. \quad (18)$$

By using the chain rule, we obtain the partial derivative of f with respect to \mathbf{W} as follows

$$\frac{\partial f}{\partial \mathbf{W}} = \frac{\partial \phi}{\partial \mathbf{W}} \frac{\partial f}{\partial \phi} = -\frac{1}{\sqrt{m}} [\mathbf{x}^T \otimes \mathbf{E}]^T \mathbf{D}^T.$$

- (ii) Let \mathbf{v} be an arbitrary vector, and \mathbf{u} be an arbitrary unit vector. The reverse triangle inequality implies that

$$\begin{aligned} \|(\mathbf{I}_m - \tilde{\gamma} \mathbf{diag}(\sigma'(\mathbf{v}))\mathbf{A})\mathbf{u}\| &\geq \|\mathbf{u}\| - \|\tilde{\gamma} \mathbf{diag}(\sigma'(\mathbf{v}))\mathbf{A}\mathbf{u}\| \\ &\geq \|\mathbf{u}\| - \tilde{\gamma} \|\mathbf{diag}(\sigma'(\mathbf{v}))\| \|\mathbf{A}\| \|\mathbf{u}\| \\ &\stackrel{(a)}{\geq} (1 - \gamma_0) \|\mathbf{u}\| \\ &= 1 - \gamma_0 > 0, \end{aligned}$$

where (a) is due to $|\sigma'(v)| \leq 1$ and $\|\mathbf{A}\|_{\text{op}} \leq c\sqrt{m}$. Therefore, taking infimum on the left-hand side over all unit vector \mathbf{u} yields the desired result.

- (iii) Since $f(\mathbf{z}^*, \mathbf{A}, \mathbf{W}) = 0$, taking implicit differentiation of f with respect to \mathbf{A} at \mathbf{z}^* gives us

$$\left(\frac{\partial \mathbf{z}}{\partial \mathbf{A}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right) \left(\frac{\partial f}{\partial \mathbf{z}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right) + \left(\frac{\partial f}{\partial \mathbf{A}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right) = 0.$$

The results in part (i)-(ii) imply the smallest eigenvalue of $\frac{\partial f}{\partial \mathbf{z}} \Big|_{\mathbf{z}=\mathbf{z}^*}$ is strictly positive, so that it is invertible. Therefore, we have

$$\frac{\partial \mathbf{z}^*}{\partial \mathbf{A}} = - \left(\frac{\partial f}{\partial \mathbf{A}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right) \left(\frac{\partial f}{\partial \mathbf{z}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right)^{-1} = \tilde{\gamma} [\mathbf{z}^T \otimes \mathbf{D}]^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}\mathbf{A}]^{-T}. \quad (19)$$

Similarly, we obtain the partial derivative of \mathbf{z}^* with respect to \mathbf{W} as follows

$$\frac{\partial \mathbf{z}^*}{\partial \mathbf{W}} = - \left(\frac{\partial f}{\partial \mathbf{W}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right) \left(\frac{\partial f}{\partial \mathbf{z}} \Big|_{\mathbf{z}=\mathbf{z}^*} \right)^{-1} = \frac{1}{\sqrt{m}} [\mathbf{x}^T \otimes \mathbf{E}]^T \mathbf{D}^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}\mathbf{A}]^{-T}. \quad (20)$$

To further simplify the notation, we denote \mathbf{z} to be the equilibrium point \mathbf{z}^* by omitting the superscribe, *i.e.*, $\mathbf{z} = \mathbf{z}^*$. Let $\hat{y} = \mathbf{u}^T \mathbf{z} + \mathbf{v}^T \phi$ be the prediction for the training data (\mathbf{x}, \mathbf{y}) . The differential of \hat{y} is given by

$$d\hat{y} = d(\mathbf{u}^T \mathbf{z} + \mathbf{v}^T \phi) = \mathbf{u}^T d\mathbf{z} + \mathbf{z} d\mathbf{u} + \mathbf{v}^T d\phi + \phi^T d\mathbf{v}.$$

The partial derivative of \hat{y} with respect to \mathbf{u} , \mathbf{v} , \mathbf{z} , and ϕ are given by

$$\frac{\partial \hat{y}}{\partial \mathbf{z}} = \mathbf{u}, \quad \frac{\partial \hat{y}}{\partial \mathbf{u}} = \mathbf{z}, \quad \frac{\partial \hat{y}}{\partial \mathbf{v}} = \phi, \quad \frac{\partial \hat{y}}{\partial \phi} = \mathbf{v}. \quad (21)$$

Let $\ell = \frac{1}{2}(\hat{y} - y)^2$. Then $\partial \ell / \partial \hat{y} = (\hat{y} - y)$. By chain rule, we have

$$\frac{\partial \ell}{\partial \mathbf{u}} = \frac{\partial \hat{y}}{\partial \mathbf{u}} \frac{\partial \ell}{\partial \hat{y}} = \mathbf{z}(\hat{y} - y), \quad (22)$$

$$\frac{\partial \ell}{\partial \phi} = \frac{\partial \hat{y}}{\partial \mathbf{v}} \frac{\partial \ell}{\partial \hat{y}} = \phi(\hat{y} - y). \quad (23)$$

By using (19)-(20) and chain rule, we obtain

$$\begin{aligned} \frac{\partial \ell}{\partial \mathbf{A}} &= \frac{\partial \mathbf{z}}{\partial \mathbf{A}} \frac{\partial \ell}{\partial \mathbf{z}} \\ &= \frac{\partial \mathbf{z}}{\partial \mathbf{A}} \frac{\partial \hat{y}}{\partial \mathbf{z}} \frac{\partial \ell}{\partial \hat{y}} = \tilde{\gamma}(\hat{y} - y) [\mathbf{z}^T \otimes \mathbf{D}]^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}\mathbf{A}]^{-T} \mathbf{u}, \end{aligned} \quad (24)$$

and

$$\begin{aligned}\frac{\partial \ell}{\partial \mathbf{W}} &= \frac{\partial \mathbf{z}}{\partial \mathbf{W}} \frac{\partial \hat{y}}{\partial \mathbf{z}} \frac{\partial \ell}{\partial \hat{y}} + \frac{\partial \phi}{\partial \mathbf{W}} \frac{\partial \hat{y}}{\partial \phi} \frac{\partial \ell}{\partial \hat{y}} \\ &= \frac{1}{\sqrt{m}} (\hat{y} - y) [\mathbf{x}^T \otimes \mathbf{E}]^T [\mathbf{D}^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D} \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}].\end{aligned}\quad (25)$$

Since $L = \sum_{i=1}^n \ell_i$ with $\ell_i = \ell(\hat{y}_i, y_i)$, we have $dL = \sum_{i=1}^n d\ell_i$ and $\partial L / \partial \ell_i = 1$. Therefore, we obtain

$$\frac{\partial L}{\partial \mathbf{A}} = \sum_{i=1}^n \frac{\partial \ell_i}{\partial \mathbf{A}} = \sum_{i=1}^n \tilde{\gamma} (\hat{y}_i - y_i) [\mathbf{z}_i^T \otimes \mathbf{D}_i]^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-T} \mathbf{u}, \quad (26)$$

$$\frac{\partial L}{\partial \mathbf{W}} = \sum_{i=1}^n \frac{\partial \ell_i}{\partial \mathbf{W}} = \sum_{i=1}^n \frac{1}{\sqrt{m}} (\hat{y}_i - y_i) [\mathbf{x}_i^T \otimes \mathbf{E}_i]^T [\mathbf{D}_i^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}], \quad (27)$$

$$\frac{\partial L}{\partial \mathbf{u}} = \sum_{i=1}^n \frac{\partial \ell_i}{\partial \mathbf{u}} = \sum_{i=1}^n (\hat{y}_i - y_i) \mathbf{z}_i, \quad (28)$$

$$\frac{\partial L}{\partial \mathbf{v}} = \sum_{i=1}^n \frac{\partial \ell_i}{\partial \mathbf{v}} = \sum_{i=1}^n (\hat{y}_i - y_i) \phi_i. \quad (29)$$

□

A.3 PROOF OF LEMMA 3.1

Proof. Let \mathbf{z}_i denote the i -th equilibrium point for the i -th data sample \mathbf{x}_i . By using (19), (20), (26) and (27), we obtain the dynamics of the equilibrium point \mathbf{z}_i as follows

$$\begin{aligned}\frac{d\mathbf{z}_i}{dt} &= \left(\frac{\partial \mathbf{z}_i}{\partial \mathbf{A}} \right)^T \frac{d\text{vec}(\mathbf{A})}{dt} + \left(\frac{\partial \mathbf{z}_i}{\partial \mathbf{W}} \right)^T \frac{d\text{vec}(\mathbf{W})}{dt} \\ &= \left(\frac{\partial \mathbf{z}_i}{\partial \mathbf{A}} \right)^T \left(-\frac{\partial L}{\partial \mathbf{A}} \right) + \left(\frac{\partial \mathbf{z}_i}{\partial \mathbf{W}} \right)^T \left(-\frac{\partial L}{\partial \mathbf{W}} \right) \\ &= -\tilde{\gamma}^2 \sum_{j=1}^n (\hat{y}_j - y_j) [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-1} [\mathbf{z}_i^T \otimes \mathbf{D}_i] [\mathbf{z}_j^T \otimes \mathbf{D}_j]^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A}]^{-T} \mathbf{u} \\ &\quad - \frac{1}{m} \sum_{j=1}^n (\hat{y}_j - y_j) [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-1} \mathbf{D}_i [\mathbf{x}_i^T \otimes \mathbf{E}_i] [\mathbf{x}_j^T \otimes \mathbf{E}_j]^T [\mathbf{D}_j^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}] \\ &= -\tilde{\gamma}^2 \sum_{j=1}^n (\hat{y}_j - y_j) [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-1} \mathbf{D}_i \mathbf{D}_j^T [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A}]^{-T} \mathbf{u} \mathbf{z}_i^T \mathbf{z}_j \\ &\quad - \frac{1}{m} \sum_{j=1}^n (\hat{y}_j - y_j) [\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A}]^{-1} \mathbf{D}_i \mathbf{E}_i \mathbf{E}_j^T [\mathbf{D}_j^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}] \mathbf{x}_i^T \mathbf{x}_j.\end{aligned}$$

By using (18) and 27, we obtain the dynamics of the feature vector ϕ_i

$$\begin{aligned}\frac{d\phi_i}{dt} &= \left(\frac{\partial \phi_i}{\partial \mathbf{W}} \right)^T \frac{d\text{vec}(\mathbf{W})}{dt} \\ &= \left(\frac{\partial \phi_i}{\partial \mathbf{W}} \right)^T \left(-\frac{\partial L}{\partial \mathbf{W}} \right) \\ &= -\frac{1}{m} \sum_{j=1}^n (\hat{y}_j - y_j) \mathbf{E}_i \mathbf{E}_j^T [\mathbf{D}_j^T (\mathbf{I}_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}] \mathbf{x}_i^T \mathbf{x}_j.\end{aligned}$$

By chain rule, the dynamics of the prediction \hat{y}_i is given by

$$\begin{aligned}
\frac{d\hat{y}_i}{dt} &= \left(\frac{\partial \hat{y}_i}{\partial \mathbf{z}_i} \right)^T \frac{d\mathbf{z}_i}{dt} + \left(\frac{\partial \hat{y}_i}{\partial \boldsymbol{\phi}_i} \right)^T \frac{d\boldsymbol{\phi}_i}{dt} + \left(\frac{\partial \hat{y}_i}{\partial \mathbf{u}} \right)^T \frac{d\mathbf{u}}{dt} + \left(\frac{\partial \hat{y}_i}{\partial \mathbf{v}} \right)^T \frac{d\mathbf{v}}{dt} \\
&= -\tilde{\gamma}^2 \sum_{j=1}^n (\hat{y}_j - y_j) \left[\mathbf{u}^T (I_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-1} \mathbf{D}_i \mathbf{D}_j^T (I_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} \right] (\mathbf{z}_i^T \mathbf{z}_j) \\
&\quad - \frac{1}{m} \sum_{j=1}^n (\hat{y}_j - y_j) \left[(\mathbf{D}_i^T (I_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-1} \mathbf{u} + \mathbf{v})^T \mathbf{E}_i \mathbf{E}_j^T (\mathbf{D}_j^T (I_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}) \right] (\mathbf{x}_i^T \mathbf{x}_j) \\
&\quad - \sum_{j=1}^n (\hat{y}_j - y_j) (\mathbf{z}_i^T \mathbf{z}_j) \\
&\quad - \sum_{j=1}^n (\hat{y}_j - y_j) (\boldsymbol{\phi}_i^T \boldsymbol{\phi}_j).
\end{aligned}$$

Define the matrices $\mathbf{M}(t) \in \mathbb{R}^{n \times n}$ and $\mathbf{Q}(t) \in \mathbb{R}^{n \times n}$ as follows

$$\begin{aligned}
\mathbf{M}(t)_{ij} &\triangleq \frac{1}{m} \mathbf{u}^T (I_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-1} \mathbf{D}_i \mathbf{D}_j^T (I_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u}, \\
\mathbf{Q}(t)_{ij} &\triangleq \frac{1}{m} (\mathbf{D}_i^T (I_m - \tilde{\gamma} \mathbf{D}_i \mathbf{A})^{-1} \mathbf{u} + \mathbf{v})^T \mathbf{E}_i \mathbf{E}_j^T (\mathbf{D}_j^T (I_m - \tilde{\gamma} \mathbf{D}_j \mathbf{A})^{-T} \mathbf{u} + \mathbf{v}).
\end{aligned}$$

Let $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\boldsymbol{\Phi}(t) \in \mathbb{R}^{n \times m}$, and $\mathbf{Z}(t) \in \mathbb{R}^{n \times m}$ be the matrices whose rows are the training data \mathbf{x}_i , feature vectors $\boldsymbol{\phi}_i$, and equilibrium points \mathbf{z}_i at time t , respectively. The dynamics of the prediction vector $\hat{\mathbf{y}}$ is given by

$$\frac{d\hat{\mathbf{y}}}{dt} = - \left[(\gamma^2 \mathbf{M}(t) + \mathbf{I}_n) \circ \mathbf{Z}(t) \mathbf{Z}(t)^T + \mathbf{Q}(t) \circ \mathbf{X} \mathbf{X}^T + \boldsymbol{\Phi}(t) \boldsymbol{\Phi}(t)^T \right] (\hat{\mathbf{y}}(t) - \mathbf{y}).$$

□

A.4 PROOF OF LEMMA 3.2

A.4.1 REVIEW OF HERMITE EXPANSIONS

To make the paper self-contained, we review the necessary background about the Hermite polynomials in this section. One can find each result in this section from any standard textbooks about functional analysis such as MacCluer (2008); Kreyszig (1978), or most recent literature (Nguyen & Mondelli, 2020, Appendix D) and (Oymak & Soltanolkotabi, 2020, Appendix H).

We consider an L^2 -space defined by $L^2(\mathbb{R}, dP)$, where dP is the *Gaussian measure*, that is,

$$dP = p(x)dx, \quad \text{where } p(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}.$$

Thus, $L^2(\mathbb{R}, dP)$ is a collection of functions f for which

$$\int_{-\infty}^{\infty} |f(x)|^2 dP(x) = \int_{-\infty}^{\infty} |f(x)|^2 p(x)dx = \mathbb{E}_{x \sim N(0,1)} |f(x)|^2 < \infty.$$

Lemma A.1. The ReLU activation $\sigma \in L^2(\mathbb{R}, dP)$.

Proof. Note that

$$\int_{-\infty}^{\infty} |\sigma(x)|^2 p(x)dx \leq \int_{-\infty}^{\infty} |x|^2 p(x)dx = \mathbb{E}_{x \sim N(0,1)} |x|^2 = \text{Var}(x) = 1.$$

□

For any functions $f, g \in L^2(\mathbb{R}, dP)$, we define an *inner product*

$$\langle f, g \rangle := \int_{-\infty}^{\infty} f(x)g(x)dP(x) = \int_{-\infty}^{\infty} f(x)g(x)p(x)dx = \mathbb{E}_{x \sim N(0,1)}[f(x)g(x)].$$

Furthermore, the induced norm $\|\cdot\|$ is given by

$$\|f\|^2 = \langle f, f \rangle = \int_{-\infty}^{\infty} |f(x)|^2 dP(x) = \mathbb{E}_{x \sim N(0,1)} |f(x)|^2.$$

This L^2 space has an orthonormal basis with respect to the inner product defined above, called *normalized probabilist's Hermite polynomials* $\{h_n(x)\}_{n=0}^{\infty}$ that are given by

$$h_n(x) = \frac{1}{\sqrt{n!}}(-1)^n e^{x^2/2} D^n(e^{-x^2/2}), \quad \text{where} \quad D^n(e^{-x^2/2}) = \frac{d^n}{dx^n} e^{-x^2/2}.$$

Lemma A.2. The *normalized probabilist's Hermite polynomials* is an orthonormal basis of $L^2(\mathbb{R}, dP)$: $\langle h_m, h_n \rangle = \delta_{mn}$.

Proof. Note that $D^n(e^{-x^2/2}) = e^{-x^2/2} P_n(x)$ for a polynomial with degree of n and leading term is $(-1)^n x^n$. Thus, we can consider $h_n(x) = \frac{1}{\sqrt{n!}}(-1)^n P_n(x)$.

Assume $m < n$

$$\begin{aligned} \langle h_n, h_m \rangle &= \mathbb{E}_{x \sim N(0,1)}[h_n(x)h_m(x)] \\ &= \int_{-\infty}^{\infty} h_n(x)h_m(x) \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx, \\ &= \frac{1}{\sqrt{2\pi}\sqrt{n!}}(-1)^n \int_{-\infty}^{\infty} D^n(e^{-x^2/2})h_m(x)dx, \quad \text{rewrite } h_n(x) \text{ by its definition} \\ &= \frac{1}{\sqrt{2\pi}\sqrt{n!}\sqrt{m!}}(-1)^{n+m} \int_{-\infty}^{\infty} D^n(e^{-x^2/2})P_m(x)dx, \quad \text{rewrite } h_m \text{ by the polynomial form} \\ &= \frac{1}{\sqrt{2\pi}\sqrt{n!}\sqrt{m!}}(-1)^{2n+m} \int_{-\infty}^{\infty} e^{-x^2/2} D_n[P_m(x)]dx, \quad \text{integration by parts } n \text{ times} \end{aligned}$$

There is no boundary terms because the super exponential decay of $e^{-x^2/2}$ at infinity. Since $m < n$, then $D_n(P_m) = 0$ so that $\langle h_m, h_n \rangle = 0$. If $m = n$, then $D_n(P_m) = (-1)^n n!$. Thus, $\langle h_n, h_n \rangle = 1$. \square

Remark: Since $\{h_n\}$ is an orthonormal basis, for every $f \in L^2(\mathbb{R}, dP)$, we have

$$f(x) = \sum_{n=0}^{\infty} \langle f, h_n \rangle h_n(x)$$

in the sense that

$$\lim_{N \rightarrow \infty} \left\| f(x) - \sum_{n=0}^N \langle f, h_n \rangle h_n(x) \right\|^2 = \lim_{N \rightarrow \infty} \mathbb{E}_{x \sim N(0,1)} \left| f(x) - \sum_{n=0}^N \langle f, h_n \rangle h_n(x) \right|^2 = 0$$

Lemma A.3. $f \in L^2(\mathbb{R}, dP)$ if and only if $\sum_{n=0}^{\infty} |\langle f, h_n \rangle|^2 < \infty$.

Proof. Note that

$$\begin{aligned}
\langle f, f \rangle &= \int_{-\infty}^{\infty} |f(x)|^2 dP(x) \\
&= \int_{-\infty}^{\infty} \left(\sum_{i=0}^{\infty} \langle f, h_i \rangle h_i(x) \right) \left(\sum_{j=0}^{\infty} \langle f, h_j \rangle h_j(x) \right) dP(x) \\
&= \sum_{i,j=0}^{\infty} \langle f, h_i \rangle \langle f, h_j \rangle \int_{-\infty}^{\infty} h_i(x) h_j(x) dP(x) \\
&= \sum_{i=1}^{\infty} |\langle f, h_i \rangle|^2.
\end{aligned}$$

□

Lemma A.4. Consider a Hilbert space H with inner product $\langle \cdot, \cdot \rangle$. If $\|f_n - f\| \rightarrow 0$ and $\|g_n - g\| \rightarrow 0$, then $\langle f, g \rangle = \lim_{n \rightarrow \infty} \langle f_n, g_n \rangle$.

Proof. Observe that

$$\begin{aligned}
|\langle f, g \rangle - \langle f_n, g_n \rangle| &\leq |\langle f, g \rangle - \langle f_n, g \rangle| + |\langle f_n, g \rangle - \langle f_n, g_n \rangle| \\
&\leq \|f\| \|g - g_n\| + \|f_n\| \|g - g_n\|.
\end{aligned}$$

Let $n \rightarrow \infty$, then the continuity of $\|\cdot\|$ implies the desired result. □

Lemma A.5. Let $\{h_n(x)\}$ be the normalized probabilist's Hermite polynomials. For any fixed number t , we have

$$e^{xt-t^2/2} = \sum_{n=0}^{\infty} \frac{t^n}{\sqrt{n!}} h_n(x). \quad (30)$$

Proof. First, we show $f(x) = e^{xt-t^2/2} \in H \triangleq L^2(\mathbb{R}, dP)$.

$$\begin{aligned}
\langle f, f \rangle &= \mathbb{E}_{x \sim N(0,1)} |f(x)|^2 \\
&= \int_{-\infty}^{\infty} e^{2xt-t^2} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\
&= e^{t^2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{(x-2t)^2}{2}\right\} dx, \quad x \sim N(2t, 1) \\
&= e^{t^2} < \infty.
\end{aligned}$$

Thus $f(x) \in H$. Then $f(x) = \sum_{n=0}^{\infty} \langle f, h_n \rangle h_n(x)$. Note that

$$\begin{aligned}
\langle f, h_n \rangle &= \mathbb{E}_{x \sim N(0,1)} [f(x) h_n(x)] \\
&= \int_{-\infty}^{\infty} e^{xt-t^2/2} \cdot \frac{1}{\sqrt{n!}} (-1)^n e^{x^2/2} D_n(e^{-x^2/2}) \cdot \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\
&= \frac{1}{\sqrt{n!}} (-1)^n \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{xt-t^2/2} \cdot D_n(e^{-x^2/2}) dx, \quad \text{integration by parts } n \text{ times} \\
&= \frac{1}{\sqrt{n!}} (-1)^{2n} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{xt-t^2/2} t^n \cdot e^{-x^2/2} dx \\
&= \frac{t^n}{\sqrt{n!}} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-(x-t)^2/2} dx, \quad x \sim N(t, 1) \\
&= \frac{t^n}{\sqrt{n!}}.
\end{aligned}$$

□

Lemma A.6. Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ with $\|\mathbf{a}\| = \|\mathbf{b}\| = 1$, then

$$\mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[h_n(\langle \mathbf{a}, \mathbf{w} \rangle)h_m(\langle \mathbf{b}, \mathbf{w} \rangle)] = \langle \mathbf{a}, \mathbf{b} \rangle^n \delta_{mn}.$$

Proof. Given fixed numbers s and t , we define two functions $f(\mathbf{w}) = e^{\langle \mathbf{a}, \mathbf{w} \rangle t - t^2/2}$ and $g(\mathbf{w}) = e^{\langle \mathbf{b}, \mathbf{w} \rangle s - s^2/2}$. Let $x = \langle \mathbf{a}, \mathbf{w} \rangle$ and $y = \langle \mathbf{b}, \mathbf{w} \rangle$. Then we have

$$\begin{aligned} f(\mathbf{w}) &= e^{\langle \mathbf{a}, \mathbf{w} \rangle t - t^2/2} = e^{xt - t^2/2} = \sum_{n=0}^{\infty} \frac{t^n}{\sqrt{n!}} h_n(x) = \sum_{n=0}^{\infty} \frac{t^n}{\sqrt{n!}} h_n(\langle \mathbf{a}, \mathbf{w} \rangle), \\ g(\mathbf{w}) &= e^{\langle \mathbf{b}, \mathbf{w} \rangle s - s^2/2} = e^{ys - s^2/2} = \sum_{n=0}^{\infty} \frac{s^n}{\sqrt{n!}} h_n(y) = \sum_{n=0}^{\infty} \frac{s^n}{\sqrt{n!}} h_n(\langle \mathbf{b}, \mathbf{w} \rangle). \end{aligned}$$

Define a Hilbert space $H_d = L^2(\mathbb{R}^d, dP)$, where dP is the *multivariate Gaussian measure*, equipped with inner product $\langle f, g \rangle \triangleq \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[f(\mathbf{w})g(\mathbf{w})]$. Clearly, $f, g \in H_d$. Define sequences $\{f_N\}$ and $\{g_N\}$ as follows

$$f_N(\mathbf{w}) = \sum_{n=0}^N \frac{t^n}{\sqrt{n!}} h_n(\langle \mathbf{a}, \mathbf{w} \rangle) \quad \text{and} \quad g_N(\mathbf{w}) = \sum_{n=0}^N \frac{s^n}{\sqrt{n!}} h_n(\langle \mathbf{b}, \mathbf{w} \rangle).$$

Since $\|f - f_N\| \rightarrow 0$ and $\|g - g_N\| \rightarrow 0$, we have

$$\begin{aligned} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[f(\mathbf{w})g(\mathbf{w})] &= \langle f, g \rangle \\ &= \lim_{N \rightarrow \infty} \langle f_N, g_N \rangle \\ &= \lim_{N \rightarrow \infty} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[f_N(\mathbf{w})g_N(\mathbf{w})] \\ &= \lim_{N \rightarrow \infty} \sum_{n,m=0}^N \frac{t^n s^m}{\sqrt{n!} \sqrt{m!}} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[h_n(\langle \mathbf{a}, \mathbf{w} \rangle)h_m(\langle \mathbf{b}, \mathbf{w} \rangle)] \end{aligned}$$

Note that the LHS is also given by

$$\begin{aligned} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[f(\mathbf{w})g(\mathbf{w})] &= e^{-t^2/2 - s^2/2} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[e^{\langle \mathbf{a}, \mathbf{w} \rangle t + \langle \mathbf{b}, \mathbf{w} \rangle s}] \\ &= e^{-t^2/2 - s^2/2} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[e^{\sum_{i=1}^d \mathbf{w}_i (a_i t + b_i s)}] \\ &= e^{-t^2/2 - s^2/2} \prod_{i=1}^d \mathbb{E}_{\mathbf{w}_i \sim N(0,1)}[e^{\mathbf{w}_i (a_i t + b_i s)}] \\ &= e^{-t^2/2 - s^2/2} \prod_{i=1}^d M_{\mathbf{w}_i}(a_i t + b_i s) \\ &= e^{\langle \mathbf{a}, \mathbf{b} \rangle st} \\ &= \sum_{n=0}^{\infty} \frac{\langle \mathbf{a}, \mathbf{b} \rangle^n (st)^n}{n!}. \end{aligned}$$

Since s and t are arbitrary numbers, matching the coefficients yields

$$\mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I}_d)}[h_n(\langle \mathbf{a}, \mathbf{w} \rangle)h_m(\langle \mathbf{b}, \mathbf{w} \rangle)] = \langle \mathbf{a}, \mathbf{b} \rangle^n \delta_{mn}.$$

□

A.4.2 LOWER BOUND THE SMALLEST EIGENVALUES OF \mathbf{G}^∞

The result in this subsection is similar to the results in (Nguyen & Mondelli, 2020, Appendix D) and (Oymak & Soltanolkotabi, 2020, Appendix H). The key difference is the assumptions made on the training data. In particular, Oymak & Soltanolkotabi (2020) assumes the training data is δ -separable, *i.e.*, $\min\{\|\mathbf{x}_i - \mathbf{x}_j\|, \|\mathbf{x}_i + \mathbf{x}_j\|\} \geq \delta > 0$ for all $i \neq j$, and Nguyen & Mondelli (2020) assumes the data \mathbf{x}_i follows some sub-Gaussian random variable, while we assume no two data are parallel to each other, *i.e.*, $\mathbf{x}_i \not\parallel \mathbf{x}_j$ for all $i \neq j$.

Lemma A.7. Given an activation function σ , if $\sigma \in L^2(\mathbb{R}, dP)$ and $\|\mathbf{x}_i\| = 1$ for all $i \in [n]$, then

$$\mathbf{G}^\infty = \sum_{k=0}^{\infty} |\langle \sigma, h_k \rangle|^2 \underbrace{(\mathbf{X}\mathbf{X}^T \circ \dots \circ \mathbf{X}\mathbf{X}^T)}_{k \text{ times}}, \quad (31)$$

where \circ is elementwise product.

Proof. Observe

$$\begin{aligned} \mathbf{G}_{ij}^\infty &= \mathbb{E}_{\mathbf{w} \sim N(0, \mathbf{I}_d)} [\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) \sigma(\langle \mathbf{w}, \mathbf{x}_j \rangle)] \\ &= \sum_{k, \ell=0}^{\infty} \langle \sigma, h_k \rangle \langle \sigma, h_\ell \rangle \mathbb{E}_{\mathbf{w} \sim N(0, \mathbf{I}_d)} [h_k(\langle \mathbf{w}, \mathbf{x}_i \rangle) h_\ell(\langle \mathbf{w}, \mathbf{x}_j \rangle)] \\ &= \sum_{k, \ell=0}^{\infty} \langle \sigma, h_k \rangle \langle \sigma, h_\ell \rangle \cdot \langle \mathbf{x}_i, \mathbf{x}_j \rangle^k \delta_{k\ell} \\ &= \sum_{k=0}^{\infty} \langle \sigma, h_k \rangle^2 \langle \mathbf{x}_i, \mathbf{x}_j \rangle^k \end{aligned}$$

□

Note that the tensor product of \mathbf{x}_i and \mathbf{x}_j is $\mathbf{x}_i \otimes \mathbf{x}_j \in \mathbb{R}^{d^2 \times 1}$, so that

$$\langle \mathbf{x}_i, \mathbf{x}_j \rangle^k = \left\langle \underbrace{\mathbf{x}_i \otimes \dots \otimes \mathbf{x}_i}_{k \text{ times}}, \underbrace{\mathbf{x}_j \otimes \dots \otimes \mathbf{x}_j}_{k \text{ times}} \right\rangle$$

Here we introduce the (*row-wise*) *Khatri–Rao product* of two matrices $\mathbf{A} \in \mathbb{R}^{k \times m}$, $\mathbf{B} \in \mathbb{R}^{k \times n}$. Then

$$\mathbf{A} * \mathbf{B} = \begin{bmatrix} \mathbf{A}_{1*} \otimes \mathbf{B}_{1*} \\ \vdots \\ \mathbf{A}_{k*} \otimes \mathbf{B}_{k*} \end{bmatrix} \in \mathbb{R}^{k \times mn}$$

where \mathbf{A}_{i*} indicates the i -th row of matrix \mathbf{A} . Therefore, the i -th row of $\mathbf{X} * \dots * \mathbf{X} \triangleq \mathbf{X}^{*n}$ is $\mathbf{x}_i \otimes \dots \otimes \mathbf{x}_i$. As a result, we obtain a more compact form of (31) as follows

$$\mathbf{G}^\infty = \sum_{k=0}^{\infty} |\langle \sigma, h_k \rangle|^2 (\mathbf{X}^{*k})(\mathbf{X}^{*k})^T. \quad (32)$$

Lemma A.8. If $\sigma(x)$ is a nonlinear function and $|\sigma(x)| \leq |x|$ and , then

$$\sup\{n : \langle \sigma, h_n \rangle > 0\} = \infty.$$

Proof. It is equivalent to show $\sigma(x)$ is not a finite linear combination of polynomials. We prove by contradiction. Suppose $\sigma(x) = a_0 + a_1x + \dots + a_nx^n$. Since $\sigma(0) = 0 = a_0$, then $\sigma(x) = a_1x + \dots + a_nx^n$. Observe that

$$\begin{aligned} \lim_{x \rightarrow \infty} \frac{|\sigma(x)|}{|x|} &= \lim_{x \rightarrow \infty} \frac{|a_1x + \dots + a_nx^n|}{|x|} \\ &= \lim_{x \rightarrow \infty} |a_1 + \dots + a_nx^{n-1}|, \\ &= \infty \end{aligned}$$

which contradicts $\frac{|\sigma(x)|}{|x|} \leq 1$ for all $x \neq 0$. □

Lemma A.9. If $\mathbf{x}_i \not\parallel \mathbf{x}_j$ for all $i \neq j$, then there exists $k_0 > 0$ such that $\lambda_{\min} [(\mathbf{X}^{*k})(\mathbf{X}^{*k})^T] > 0$ for all $k \geq k_0$. Therefore, $\lambda_{\min}(\mathbf{G}^\infty) > 0$.

Proof. To simplify the notation, denote $\mathbf{K} = (\mathbf{X}^{*k})^T \in \mathbb{R}^{kd \times n}$. Since $\mathbf{x}_i \perp \mathbf{x}_j$ and $\|\mathbf{x}_i\| = 1$, then let $\delta \triangleq \max\{|\langle \mathbf{x}_i, \mathbf{x}_j \rangle|\} = \max\{|\cos \theta_{ij}|\}$ and $\delta \in (0, 1)$, where θ_{ij} is the angle between \mathbf{x}_i and \mathbf{x}_j . For any unit vector $\mathbf{v} \in \mathbb{R}^n$, we have

$$\begin{aligned} \mathbf{v}^T (\mathbf{X}^{*k}) (\mathbf{X}^{*k})^T \mathbf{v} &= \|\mathbf{K}\mathbf{v}\|^2 = \left\| \sum_{i=1}^n v_i \mathbf{K}_{*i} \right\|^2 \\ &= \sum_{i=1}^n \sum_{j=1}^n v_i v_j \langle \mathbf{K}_{*i}, \mathbf{K}_{*j} \rangle \\ &= \sum_{i=1}^n \sum_{j=1}^n v_i v_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle^k \\ &= \sum_{i=1}^n v_i^2 \|\mathbf{x}_i\|^{2k} + \sum_{i \neq j} v_i v_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle^k \\ &= 1 + \sum_{i \neq j} v_i v_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle^k, \end{aligned}$$

where the last equality is because $\|\mathbf{x}_i\| = 1$ and $\|\mathbf{v}\| = 1$. Note that

$$\begin{aligned} \left| \sum_{i \neq j} v_i v_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle^k \right| &\leq \sum_{i \neq j} |v_i| |v_j| |\langle \mathbf{x}_i, \mathbf{x}_j \rangle|^k \\ &\leq \delta^k \sum_{i \neq j} |v_i| |v_j|, \quad \text{by } |\langle \mathbf{x}_i, \mathbf{x}_j \rangle| \leq \delta \\ &\leq \delta^k \left(\sum_{i=1}^n |v_i| \right)^2 \\ &\leq n \delta^k, \quad \text{by Cauchy-Schwarz's inequality.} \end{aligned}$$

By inverse triangle inequality, we have

$$\|\mathbf{K}\mathbf{v}\|^2 \geq 1 - n\delta^k.$$

Choose $k_0 \geq \log n / \log(1/\delta)$, then $\lambda_{\min}\{(\mathbf{X}^{*k})(\mathbf{X}^{*k})^T\} > 0$ for all $k \geq k_0$. \square

A.5 PROOF OF LEMMA 3.3

Proof. Since $\|\mathbf{x}_i\| = 1$ and $\mathbf{w}_r(0) \sim \mathcal{N}(0, \mathbf{I}_d)$, we have $\mathbf{x}_i^T \mathbf{w}_r(0) \sim N(0, 1)$ for all $i \in [n]$. Let $X_{ir} \triangleq \sigma [\mathbf{x}_i^T \mathbf{w}_r(0)]$ and $Z \sim N(0, 1)$, then for any $|\lambda| \leq 1/\sqrt{2}$, we have

$$\mathbb{E} \exp\{X_{ir}^2 \lambda^2\} = \mathbb{E} \exp\{\sigma [\mathbf{x}_i^T \mathbf{w}_r(0)]^2 \lambda^2\} \leq \mathbb{E} \exp\{Z^2 \lambda^2\} = 1/\sqrt{1-2\lambda^2} \leq e^{2t^2},$$

where the first inequality is due to $|\sigma(x)| \leq |x|$, and the last inequality is by using the numerical inequality $1/(1-x) \leq e^{2x}$. Choose $\lambda \leq (\log \sqrt{2})^{1/2}$, we obtain $\mathbb{E}\{X_{ir}^2 \lambda^2\} \leq 2$. By using Markov's inequality, we have for any $t \geq 0$

$$\mathbb{P}\{|X_{ir}| \geq t\} = \mathbb{P}\{|X_{ir}|^2 / \log \sqrt{2} \geq t^2 / \log \sqrt{2}\} \leq 2 \exp\{-t^2 \log \sqrt{2}\} \leq 2 \exp\{-t^2/4\}.$$

Therefore X_{ir} is a sub-Gaussian random variable with sub-Gaussian norm $\|X_{ir}\|_{\psi_2} \leq 2$ (Vershynin, 2018, Proposition 2.5.2). Then $X_{ir} X_{jr}$ is a sub-exponential random variable with sub-exponential norm $\|X_{ir} X_{jr}\|_{\psi_1} \leq 4$ (Vershynin, 2018, Lemma 2.7.7). Observe that

$$\mathbf{G}_{ij}(0) = \phi_i(0)^T \phi_j(0) = \frac{1}{m} \sum_{r=1}^m \sigma [\mathbf{x}_i^T \mathbf{w}_r(0)] \sigma [\mathbf{x}_j^T \mathbf{w}_r(0)] = \frac{1}{m} \sum_{r=1}^m X_{ir} X_{jr}.$$

Since $\mathbf{G}_{ij}^\infty = \mathbb{E}[\mathbf{G}_{ij}(0)]$, (Vershynin, 2018, Exercise 2.7.10) implies that $\mathbf{G}_{ij}(0) - \mathbf{G}_{ij}^\infty$ is also a zero-mean sub-exponential random variable. It follows from the Bernstein's inequality that

$$\begin{aligned}
\mathbb{P}\left\{\|\mathbf{G}(0) - \mathbf{G}^\infty\|_2 \geq \frac{\lambda_0}{4}\right\} &\leq \mathbb{P}\left\{\|\mathbf{G}(0) - \mathbf{G}^\infty\|_F \geq \frac{\lambda_0}{4}\right\} \\
&= \mathbb{P}\left\{\|\mathbf{G}(0) - \mathbf{G}^\infty\|_F^2 \geq \left(\frac{\lambda_0}{4}\right)^2\right\} \\
&= \mathbb{P}\left\{\sum_{i,j=1}^n |\mathbf{G}_{ij}(0) - \mathbf{G}_{ij}^\infty|^2 \geq \left(\frac{\lambda_0}{4}\right)^2\right\} \\
&\leq \sum_{i,j=1}^n \mathbb{P}\left\{|\mathbf{G}_{ij}(0) - \mathbf{G}_{ij}^\infty|^2 \geq \left(\frac{\lambda_0}{4n}\right)^2\right\} \\
&= \sum_{i,j=1}^n \mathbb{P}\left\{|\mathbf{G}_{ij}(0) - \mathbf{G}_{ij}^\infty| \geq \frac{\lambda_0}{4n}\right\} \\
&\leq n^2 \cdot 2 \exp\{-c\lambda_0^2 m/n^2\} \\
&\leq \delta,
\end{aligned}$$

where $c > 0$ is some constant, and we use the facts $\|\mathbf{X}\|_2 \leq \|\mathbf{X}\|_F$, and $\mathbb{P}\{\sum_{i=1}^n x_i \geq \varepsilon\} \leq \sum_{i=1}^n \mathbb{P}\{x_i \geq \varepsilon/n\}$. \square

A.6 PROOF OF LEMMA 3.4

Proof. By using the 1-Lipschitz continuity of $\sigma(x)$, we have

$$\begin{aligned}
\|\mathbf{G} - \mathbf{G}(0)\| &= \frac{1}{m} \|\sigma(\mathbf{X}\mathbf{W}^T)\sigma(\mathbf{X}\mathbf{W}^T)^T - \sigma(\mathbf{X}\mathbf{W}(0)^T)\sigma(\mathbf{X}\mathbf{W}(0)^T)^T\| \\
&\leq \frac{1}{m} \|\sigma(\mathbf{X}\mathbf{W}^T)\sigma(\mathbf{X}\mathbf{W}^T)^T - \sigma(\mathbf{X}\mathbf{W}^T)\sigma(\mathbf{X}\mathbf{W}(0)^T)^T\| \\
&\quad + \frac{1}{m} \|\sigma(\mathbf{X}\mathbf{W}^T)\sigma(\mathbf{X}\mathbf{W}(0)^T)^T - \sigma(\mathbf{X}\mathbf{W}(0)^T)\sigma(\mathbf{X}\mathbf{W}(0)^T)^T\| \\
&= \frac{1}{m} \|\sigma(\mathbf{X}\mathbf{W}^T)\| \|\sigma(\mathbf{X}\mathbf{W}^T) - \sigma(\mathbf{X}\mathbf{W}(0)^T)\| \\
&\quad + \frac{1}{m} \|\sigma(\mathbf{X}\mathbf{W}^T) - \sigma(\mathbf{X}\mathbf{W}(0)^T)\| \|\sigma(\mathbf{X}\mathbf{W}(0)^T)\| \\
&\leq \frac{1}{m} \|\mathbf{X}\| \|\mathbf{W}\| \|\mathbf{X}\| \|\mathbf{W} - \mathbf{W}(0)\| + \frac{1}{m} \|\mathbf{X}\| \|\mathbf{W} - \mathbf{W}(0)\| \|\mathbf{X}\| \|\mathbf{W}(0)\| \\
&\leq \frac{4c}{\sqrt{m}} \|\mathbf{X}\|^2 \|\mathbf{W} - \mathbf{W}(0)\| \\
&\leq \frac{\lambda_0}{4}.
\end{aligned}$$

\square

A.7 PROOF OF LEMMA 3.5

Proof. It suffices to show the result holds for $\gamma = \min\{\gamma_0, \gamma_0/2c\}$, where $\gamma_0 = 1/2$. Note that Lemma 2.1 still holds if one chooses a larger c . Thus, we choose $c \gtrsim \sqrt{\lambda_0}/\|\mathbf{X}\|$. We prove by the induction. Suppose that for $0 \leq s \leq t$, the followings hold

- (i) $\lambda_{\min}(\mathbf{G}(s)) \geq \frac{\lambda_0}{2}$,
- (ii) $\|\mathbf{u}(s)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iii) $\|\mathbf{v}(s)\| \leq \frac{8c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,

-
- (iv) $\|\mathbf{W}(s)\| \leq 2c\sqrt{m}$,
 - (v) $\|\mathbf{A}(s)\| \leq 2c\sqrt{m}$,
 - (vi) $\|\hat{\mathbf{y}}(s) - \mathbf{y}\|^2 \leq \exp\{-\lambda_0 s\} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2$,

Since $\lambda_{\min}(\mathbf{G}(s)) \geq \frac{\lambda_0}{2}$, we have

$$\begin{aligned} \frac{d}{dt} \|\hat{\mathbf{y}}(t) - \mathbf{y}\|^2 &= -2(\hat{\mathbf{y}}(t) - \mathbf{y})^T \mathbf{H}(t)(\hat{\mathbf{y}}(t) - \mathbf{y}) \\ &\leq -\lambda_0 \|\hat{\mathbf{y}}(t) - \mathbf{y}\|^2 \end{aligned}$$

Solving the ordinary differential equation yields

$$\|\hat{\mathbf{y}}(t) - \mathbf{y}\|^2 \leq \exp\{-\lambda_0 t\} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2.$$

By using the inductive hypothesis $\|\mathbf{W}(s)\| \leq 2c\sqrt{m}$, we have

$$\|\phi_i(s)\| = \left\| \frac{1}{\sqrt{m}} \sigma(\mathbf{W}(s)\mathbf{x}_i) \right\| \leq \frac{1}{\sqrt{m}} \|\mathbf{W}(s)\| \|\mathbf{x}_i\| \leq 2c.$$

It follows from Lemma 2.2 with $\gamma_0 = 1/2$ that

$$\|\mathbf{z}_i^*(s)\| \leq 2\|\phi_i(s)\| \leq 4c.$$

Note that

$$\begin{aligned} \|\nabla_{\mathbf{v}} L(s)\| &\leq \sum_{i=1}^n |\hat{y}_i(s) - y_i| \|\phi_i(s)\| \\ &\leq 2c \sum_{i=1}^n |\hat{y}_i(s) - y_i| \\ &\leq 2c\sqrt{n} \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\ &\leq 2c\sqrt{n} \exp\{-\lambda_0 s/2\} \|\mathbf{y}(0) - \mathbf{y}\| \end{aligned}$$

and so

$$\begin{aligned} \|\mathbf{v}(t) - \mathbf{v}(0)\| &\leq \int_0^t \|\nabla_{\mathbf{v}} L(s)\| ds \\ &\leq 2c\sqrt{n} \|\mathbf{y}(0) - \mathbf{y}\| \int_0^t \exp\{-\lambda_0 s/2\} ds \\ &\leq \frac{4c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|, \end{aligned}$$

Since $\mathbf{v}_i(0)$ follows symmetric Bernoulli distribution, then $\|\mathbf{v}(0)\| = \sqrt{m}$ and we obtain

$$\|\mathbf{v}(t)\| \leq \|\mathbf{v}(t) - \mathbf{v}(0)\| + \|\mathbf{v}(0)\| \leq \frac{8c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|,$$

where the last inequality is due to $m = \Omega\left(\frac{c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2\right)$ and $c \gtrsim \sqrt{\lambda_0} / \|\mathbf{X}\|$.

Similarly, we have

$$\begin{aligned} \|\nabla_{\mathbf{u}} L(s)\| &\leq \sum_{i=1}^n |\hat{y}_i(s) - y_i| \|\mathbf{z}_i^*\| \\ &\leq 4c\sqrt{n} \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\ &\leq 4c\sqrt{n} \exp\{-\lambda_0 s/2\} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|, \end{aligned}$$

so that

$$\|\mathbf{u}(t) - \mathbf{u}(0)\| \leq \int_0^t \|\nabla_{\mathbf{u}} L(s)\| ds \leq \frac{8c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|.$$

Since $\mathbf{u}_i(0)$ follows symmetric Bernoulli distribution, then $\|\mathbf{u}(0)\| = \sqrt{m}$ and we obtain

$$\|\mathbf{u}(t)\| \leq \|\mathbf{u}(t) - \mathbf{u}(0)\| + \|\mathbf{u}(0)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|.$$

Note that

$$\begin{aligned} \|\nabla_{\mathbf{W}} L(s)\| &\leq \sum_{i=1}^n \frac{1}{\sqrt{m}} |\hat{y}_i(s) - y_i| \|\mathbf{E}_i(s)\| (\|\mathbf{U}_i(s)^{-1} \mathbf{u}(s)\| + \|\mathbf{v}(s)\|) \|\mathbf{x}_i\| \\ &\leq \frac{64c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \sum_{i=1}^n |\hat{y}_i(s) - y_i| \\ &\leq \frac{64cn}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\ &\leq \frac{64cn}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \exp\{-\lambda_0 s/2\}, \end{aligned}$$

so that

$$\begin{aligned} \|\mathbf{W}(t) - \mathbf{W}(0)\| &\leq \int_0^t \|\nabla_{\mathbf{W}} L(s)\| ds \\ &\leq \frac{128cn}{\lambda_0^2 \sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \\ &\leq \frac{\lambda_0 \sqrt{m}}{16c \|\mathbf{X}\|^2} \\ &\leq R. \end{aligned}$$

Therefore, we obtain

$$\|\mathbf{W}(t)\| \leq \|\mathbf{W}(t) - \mathbf{W}(0)\| + \|\mathbf{W}(0)\| \leq 2c\sqrt{m},$$

Moreover, it follows from Lemma 3.4 that $\lambda_{\min}\{\mathbf{G}(t)\} \geq \frac{\lambda_0}{2}$.

Note that

$$\begin{aligned} \|\nabla_{\mathbf{A}} L(s)\| &\leq \sum_{i=1}^n \frac{\gamma}{\sqrt{m}} |\hat{y}_i(s) - y_i| \|\mathbf{D}_i\| \|\mathbf{U}_i(s)^{-1}\| \|\mathbf{u}(s)\| \|\mathbf{z}_i^*\| \\ &\leq \frac{32c\sqrt{n}}{\lambda_0 \sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \sum_{i=1}^n |\hat{y}_i(s) - y_i| \\ &\leq \frac{32cn}{\lambda_0 \sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\ &\leq \frac{32cn}{\lambda_0 \sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \exp\{-\lambda_0 s/2\}, \end{aligned}$$

so that

$$\begin{aligned} \|\mathbf{A}(t) - \mathbf{A}(0)\| &\leq \int_0^t \|\nabla_{\mathbf{A}} L(s)\| ds \\ &\leq \frac{64cn}{\lambda_0^2 \sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2. \end{aligned}$$

Then

$$\|\mathbf{A}(t)\| \leq \|\mathbf{A}(t) - \mathbf{A}(0)\| + \|\mathbf{A}(0)\| \leq 2c\sqrt{m}.$$

□

A.8 PROOF OF LEMMA 3.6

In this section, we prove the result for discrete time analysis or result for gradient descent. Assume $\|\mathbf{A}(0)\| \leq c\sqrt{m}$ and $\|\mathbf{W}(0)\| \leq c\sqrt{m}$. Further, we assume $\lambda_{\min}(\mathbf{G}(0)) \geq \frac{3}{4}\lambda_0$ and we assume $m = \Omega\left(\frac{c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2\right)$ and choose $0 < \gamma \leq \min\{1/2, 1/4c\}$. Moreover, we assume the stepsize $\alpha = \mathcal{O}(\lambda_0/n^2)$. We make the inductive hypothesis as follows for all $0 \leq s \leq k$

- (i) $\lambda_{\min}(\mathbf{G}(s)) \geq \frac{\lambda_0}{2}$,
- (ii) $\|\mathbf{u}(s)\| \leq \frac{32c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iii) $\|\mathbf{v}(s)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|$,
- (iv) $\|\mathbf{W}(s)\| \leq 2c\sqrt{m}$,
- (v) $\|\mathbf{A}(s)\| \leq 2c\sqrt{m}$,
- (vi) $\|\hat{\mathbf{y}}(s) - \mathbf{y}\|^2 \leq (1 - \alpha\lambda_0/2)^s \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2$.

Proof. Note that Lemma 2.1 still holds if one chooses a larger c . Thus, we choose $c \gtrsim \sqrt{\lambda_0}/\|\mathbf{X}\|$. By using the inductive hypothesis, we have for any $0 \leq s \leq k$

$$\|\phi_i(s)\| = \left\| \frac{1}{\sqrt{m}} \sigma(\mathbf{W}(s)\mathbf{x}_i) \right\| \leq \frac{1}{\sqrt{m}} \|\mathbf{W}(s)\| \leq 2c$$

and

$$\|\Phi(s)\| \leq \|\Phi(s)\|_F = \left(\sum_{i=1}^n \|\phi_i(s)\|^2 \right)^{1/2} \leq 2c\sqrt{n}. \quad (33)$$

By using Lemma 2.2, we obtain the upper bound for the equilibrium point $\mathbf{z}_i(s)$ for any $0 \leq s \leq k$ as follows

$$\|\mathbf{z}_i(s)\| \leq \frac{1}{1 - \gamma_0} \|\phi_i(s)\| = 2\|\phi_i(s)\| \leq 4c,$$

where the last inequality is because we choose $\gamma_0 = 1/2$, and

$$\|\mathbf{Z}(s)\| \leq \|\mathbf{Z}(s)\|_F = \left(\sum_{i=1}^n \|\mathbf{z}_i(s)\|^2 \right)^{1/2} = 4c\sqrt{n}. \quad (34)$$

By using the upper bound of $\phi_i(s)$, we obtain for any $0 \leq s \leq k$

$$\begin{aligned} \|\nabla_{\mathbf{v}} L(s)\| &\leq \sum_{i=1}^n |\hat{y}_i(s) - y_i| \|\phi_i(s)\| \\ &\leq 2c \sum_{i=1}^n |\hat{y}_i(s) - y_i| \\ &\leq 2c\sqrt{n} \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\ &\leq 2c\sqrt{n} (1 - \alpha\lambda_0/2)^{s/2} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|. \end{aligned}$$

Let $\beta \triangleq \sqrt{1 - \alpha\lambda_0/2}$. Then the upper bound of $\|\nabla_{\mathbf{v}} L(s)\|$ can be written as

$$\|\nabla_{\mathbf{v}} L(s)\| \leq 2c\sqrt{n}\beta^s \|\hat{\mathbf{y}}(0) - \mathbf{y}\|, \quad (35)$$

and

$$\begin{aligned}
\|\mathbf{v}(k+1) - \mathbf{v}(0)\| &\leq \sum_{s=0}^k \|\mathbf{v}(s+1) - \mathbf{v}(s)\| = \alpha \sum_{s=0}^k \|\nabla_{\mathbf{v}} L(s)\| \\
&\leq \alpha \cdot 2c\sqrt{n} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \sum_{s=0}^k \beta^s \\
&= \frac{2(1-\beta^2)}{\lambda_0} \cdot 2c\sqrt{n} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \frac{1-\beta^{k+1}}{1-\beta} \\
&\leq \frac{8c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|,
\end{aligned}$$

where the last inequality we use the facts $\beta < 1$. By triangle inequality, we obtain

$$\|\mathbf{v}(k+1)\| \leq \|\mathbf{v}(k+1) - \mathbf{v}(0)\| + \|\mathbf{v}(0)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|,$$

which proves the result (iii). Similarly, we can upper bound the gradient of \mathbf{u}

$$\|\nabla_{\mathbf{u}} L(s)\| \leq \sum_{i=1}^n |\hat{y}_i(s) - y_i| \|\mathbf{z}_i\| \leq 4c\sqrt{n} \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \leq 4c\sqrt{n}\beta^s \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \quad (36)$$

so that

$$\|\mathbf{u}(k+1) - \mathbf{u}(0)\| \leq \frac{16c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}} - \mathbf{y}\|,$$

and

$$\|\mathbf{u}(k)\| \leq \|\mathbf{u}(k) - \mathbf{u}(0)\| + \|\mathbf{u}(0)\| \leq \frac{32c\sqrt{n}}{\lambda_0} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|.$$

The result (ii) is also obtained.

By using the inductive hypothesis, we can upper bound the gradient of \mathbf{W} as follows

$$\begin{aligned}
\|\nabla_{\mathbf{W}} L(s)\| &\leq \sum_{i=1}^n \frac{1}{\sqrt{m}} |\hat{y}_i(s) - y_i| \|\mathbf{E}_i(s)\| (\|\mathbf{U}_i(s)^{-1} \mathbf{u}(s)\| + \|\mathbf{v}(s)\|) \|\mathbf{x}_i\| \\
&\leq \frac{128c\sqrt{n}}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \sum_{i=1}^n |\hat{y}_i(s) - y_i| \\
&\leq \frac{128cn}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\
&\leq \frac{128cn}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^s,
\end{aligned} \quad (37)$$

so that

$$\begin{aligned}
\|\mathbf{W}(k+1) - \mathbf{W}(0)\| &\leq \alpha \sum_{s=0}^k \|\nabla_{\mathbf{W}} L(s)\| \\
&\leq \alpha \cdot \frac{128cn}{\lambda_0^2\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \sum_{s=0}^k \beta^s \\
&\leq \frac{512cn}{\lambda_0^2\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \\
&\leq \frac{\sqrt{m}\lambda_0}{16c\|\mathbf{X}\|^2} \\
&\leq R,
\end{aligned}$$

where the third inequality holds is because m is large, *i.e.*, $m = \Theta\left(\frac{c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2\right)$. To simplify the notation, we assume

$$m = \frac{C c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \quad (38)$$

for some large number $C > 0$. Moreover, we obtain

$$\|\mathbf{W}(k+1)\| \leq \|\mathbf{W}(k+1) - \mathbf{W}(0)\| + \|\mathbf{W}(0)\| \leq 2c\sqrt{m},$$

Therefore, it follows from Lemma 3.4 that $\lambda_{\min}\{\mathbf{G}(k+1)\} \geq \frac{\lambda_0}{2}$. Thus, the results (i) and (iv) are established.

By using similar argument, we can upper bound the gradient of \mathbf{A} as follows Note that

$$\begin{aligned} \|\nabla_{\mathbf{A}} L(s)\| &\leq \sum_{i=1}^n \frac{\gamma}{\sqrt{m}} |\hat{y}_i(s) - y_i| \|\mathbf{D}_i\| \|\mathbf{U}_i(s)^{-1}\| \|\mathbf{u}(s)\| \|\mathbf{z}_i^*\| \\ &\leq \frac{64c\sqrt{n}}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \sum_{i=1}^n |\hat{y}_i(s) - y_i| \\ &\leq \frac{64cn}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \|\hat{\mathbf{y}}(s) - \mathbf{y}\| \\ &\leq \frac{64cn}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^s, \end{aligned}$$

so that

$$\begin{aligned} \|\mathbf{A}(k+1) - \mathbf{A}(0)\| &\leq \alpha \sum_{s=0}^k \|\nabla_{\mathbf{A}} L(s)\| \\ &\leq \alpha \cdot \frac{64cn}{\lambda_0\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \sum_{s=0}^k \beta^s \\ &\leq \frac{256cn}{\lambda_0^2\sqrt{m}} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2. \end{aligned}$$

Since $m = \frac{C c^2 n \|\mathbf{X}\|^2}{\lambda_0^3} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2$ and $c, C > 0$ are large enough, we have

$$\|\mathbf{A}(k+1)\| \leq \|\mathbf{A}(k+1) - \mathbf{A}(0)\| + \|\mathbf{A}(0)\| \leq 2c\sqrt{m}.$$

Therefore, the result (v) is obtained and the equilibrium points $\mathbf{z}_i(k+1)$ exists for all $i \in [n]$.

To establish the result (vi), we need to derive the bounds between equilibrium points $\mathbf{Z}(k)$ and feature vectors $\Phi(k)$. We firstly bound the difference between equilibrium points $\mathbf{z}_i(k+1)$ and $\mathbf{z}_i(k)$. For any $\ell \geq 1$, we have

$$\begin{aligned} \|\mathbf{z}_i^{\ell+1}(k+1) - \mathbf{z}_i^{\ell+1}(k)\| &= \|\sigma [\tilde{\gamma} \mathbf{A}(k+1) \mathbf{z}_i^\ell(k+1) + \phi_i(k+1)] - \sigma [\tilde{\gamma} \mathbf{A}(k) \mathbf{z}_i^\ell(k) + \phi_i(k)]\| \\ &\leq \|\tilde{\gamma} \mathbf{A}(k+1) \mathbf{z}_i^\ell(k+1) + \phi_i(k+1) - \tilde{\gamma} \mathbf{A}(k) \mathbf{z}_i^\ell(k) - \phi_i(k)\| \\ &\leq \tilde{\gamma} \|\mathbf{A}(k+1) \mathbf{z}_i^\ell(k+1) - \mathbf{A}(k) \mathbf{z}_i^\ell(k)\| + \|\phi_i(k+1) - \phi_i(k)\|, \end{aligned}$$

where the first term can be bounded as follows

$$\begin{aligned} &\tilde{\gamma} \|\mathbf{A}(k+1) \mathbf{z}_i^\ell(k+1) - \mathbf{A}(k) \mathbf{z}_i^\ell(k)\| \\ &\leq \tilde{\gamma} \|\mathbf{A}(k+1) - \mathbf{A}(k)\| \|\mathbf{z}_i^\ell(k+1)\| + \tilde{\gamma} \|\mathbf{A}(k)\| \|\mathbf{z}_i^\ell(k+1) - \mathbf{z}_i^\ell(k)\| \\ &\leq \tilde{\gamma} \alpha \|\nabla_{\mathbf{A}} L(k)\| (4c) + \tilde{\gamma} \|\mathbf{A}(k)\| \|\mathbf{z}_i^\ell(k+1) - \mathbf{z}_i^\ell(k)\| \\ &\leq \frac{64\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \beta^k + (1/2) \|\mathbf{z}_i^\ell(k+1) - \mathbf{z}_i^\ell(k)\|, \end{aligned}$$

and the second term is bounded as follows

$$\begin{aligned}
& \|\phi_i(k+1) - \phi_i(k)\| \\
&= \frac{1}{\sqrt{m}} \|\sigma[\mathbf{W}(k+1)\mathbf{x}_i] - \sigma[\mathbf{W}(k)\mathbf{x}_i]\| \\
&\leq \frac{1}{\sqrt{m}} \|\mathbf{W}(k+1) - \mathbf{W}(k)\| \|\mathbf{x}_i\| \\
&\leq \frac{\alpha}{\sqrt{m}} \|\nabla_{\mathbf{W}} L(k)\| \\
&\leq \frac{128\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k.
\end{aligned}$$

Thus, we obtain

$$\begin{aligned}
\|\mathbf{z}_i^{\ell+1}(k+1) - \mathbf{z}_i^{\ell+1}(k)\| &\leq (1/2) \|\mathbf{z}_i^{\ell}(k+1) - \mathbf{z}_i^{\ell}(k)\| + \frac{256\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k \\
&\leq (1/2)^\ell \|\mathbf{z}_i^1(k+1) - \mathbf{z}_i^1(k)\| + \frac{256\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k \cdot \sum_{j=0}^{\ell-1} 2^{-j} \\
&\leq (1/2)^\ell \|\mathbf{z}_i^1(k+1) - \mathbf{z}_i^1(k)\| + \frac{512\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k.
\end{aligned}$$

By letting $\ell \rightarrow \infty$, we obtain

$$\|\mathbf{z}_i(k+1) - \mathbf{z}_i(k)\| \leq \frac{512\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k.$$

By using the Cauchy-Schwartz's inequality, we have

$$\|\mathbf{Z}(k+1) - \mathbf{Z}(k)\| \leq \|\mathbf{Z}(k+1) - \mathbf{Z}(k)\|_F \leq \frac{512\alpha cn^{3/2}}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k. \quad (39)$$

In addition, we will also bound the difference in $\phi_i(k+1)$ and $\phi_i(k)$. Note that

$$\|\phi_i(k+1) - \phi_i(k)\| = \frac{1}{\sqrt{m}} \|\sigma[\mathbf{W}(k+1)\mathbf{x}_i] - \sigma[\mathbf{W}(k)\mathbf{x}_i]\| \leq \frac{128\alpha cn}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k,$$

so that

$$\|\Phi(k+1) - \Phi(k)\| \leq \|\Phi(k+1) - \Phi(k)\|_F \leq \frac{128\alpha cn^{3/2}}{\lambda_0 m} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \cdot \beta^k \quad (40)$$

Now, we are ready to establish the result (vi). Note that

$$\begin{aligned}
\|\hat{\mathbf{y}}(k+1) - \mathbf{y}\|^2 &= \|\hat{\mathbf{y}}(k+1) - \hat{\mathbf{y}}(k) + \hat{\mathbf{y}}(k) - \mathbf{y}\|^2 \\
&= \|\hat{\mathbf{y}}(k+1) - \hat{\mathbf{y}}(k)\|^2 + 2 \langle \hat{\mathbf{y}}(k+1) - \hat{\mathbf{y}}(k), \hat{\mathbf{y}}(k) - \mathbf{y} \rangle + \|\hat{\mathbf{y}}(k) - \mathbf{y}\|^2.
\end{aligned}$$

In the rest of this proof, we will bound each term in the above inequality. By the prediction rule of $\hat{\mathbf{y}}$, we can bound the difference between $\hat{\mathbf{y}}(k+1)$ and $\hat{\mathbf{y}}(k)$ as follows

$$\begin{aligned}
\|\hat{\mathbf{y}}(k+1) - \hat{\mathbf{y}}(k)\| &= \|\mathbf{Z}(k+1)\mathbf{u}(k+1) + \Phi(k+1)\mathbf{v}(k+1) - \mathbf{Z}(k)\mathbf{u}(k) - \Phi(k)\mathbf{v}(k)\| \\
&\leq \|\mathbf{Z}(k+1)\mathbf{u}(k+1) - \mathbf{Z}(k)\mathbf{u}(k)\| + \|\Phi(k+1)\mathbf{v}(k+1) - \Phi(k)\mathbf{v}(k)\|,
\end{aligned}$$

where the first term can be bounded as follows by using (34), 36, 38, 39, hypothesis (ii), and a large constant $C_0 > 0$

$$\begin{aligned}
& \|\mathbf{Z}(k+1)\mathbf{u}(k+1) - \mathbf{Z}(k)\mathbf{u}(k)\| + \|\mathbf{Z}(k+1) - \mathbf{Z}(k)\| \|\mathbf{u}(k)\| \\
&= \alpha \|\mathbf{Z}(k+1)\| \|\nabla_{\mathbf{u}} L(k)\| + \|\mathbf{Z}(k+1) - \mathbf{Z}(k)\| \|\mathbf{u}(k)\| \\
&\leq \alpha C_0 c^2 n \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \beta^k,
\end{aligned}$$

and the second term is bounded as follows by using (33), 35, 40, 38, hypothesis (iii), and a large constant $C_0 > 0$

$$\begin{aligned} & \|\Phi(k+1)\|\|\mathbf{v}(k+1) - \mathbf{v}(k)\| + \|\Phi(k+1) - \Phi(k)\|\|\mathbf{v}(k)\| \\ &= \alpha\|\Phi(k+1)\|\|\nabla_{\mathbf{v}}L(k)\| + \|\Phi(k+1) - \Phi(k)\|\|\mathbf{v}(k)\| \\ &\leq \alpha C_0 c^2 n \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \beta^k. \end{aligned}$$

Therefore, we have

$$\|\hat{\mathbf{y}}(k+1) - \hat{\mathbf{y}}(k)\| \leq \alpha C_0 c^2 n \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \beta^k, \quad (41)$$

where the scalar 2 is absorbed in C_0 and the constant C_0 is difference from C .

Let $\mathbf{g} \triangleq \mathbf{Z}(k)\mathbf{u}(k+1) + \Phi(k)\mathbf{v}(k+1)$. Then we have

$$\langle \hat{\mathbf{y}}(k+1) - \hat{\mathbf{y}}(k), \hat{\mathbf{y}}(k) - \mathbf{y} \rangle = \langle \hat{\mathbf{y}}(k+1) - \mathbf{g}, \hat{\mathbf{y}}(k) - \mathbf{y} \rangle + \langle \mathbf{g} - \hat{\mathbf{y}}(k), \hat{\mathbf{y}}(k) - \mathbf{y} \rangle.$$

Let us bound each term individually. By using the Cauchy-Schwartz inequality, we have

$$\begin{aligned} & \langle \hat{\mathbf{y}}(k+1) - \mathbf{g}, \hat{\mathbf{y}}(k) - \mathbf{y} \rangle \\ &= \langle (\mathbf{Z}(k+1) - \mathbf{Z}(k))\mathbf{u}(k+1), \hat{\mathbf{y}}(k) - \mathbf{y} \rangle + \langle (\Phi(k+1) - \Phi(k))\mathbf{v}(k+1), \hat{\mathbf{y}}(k) - \mathbf{y} \rangle \\ &\leq (\|\mathbf{Z}(k+1) - \mathbf{Z}(k)\|\|\mathbf{u}(k+1)\| + \|\Phi(k+1) - \Phi(k)\|\|\mathbf{v}(k+1)\|) \|\hat{\mathbf{y}}(k) - \mathbf{y}\| \\ &\leq \alpha C_0 c^2 n \|\hat{\mathbf{y}}(0) - \mathbf{y}\| \cdot \beta^k \|\hat{\mathbf{y}}(k) - \mathbf{y}\|, \quad \text{by (34), 36, 38, 39} \\ &\leq \alpha C_0 c^2 n \cdot \beta^{2k} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2. \end{aligned} \quad (42)$$

By using $\nabla_{\mathbf{u}}L(k) = \mathbf{Z}(k)^T(\hat{\mathbf{y}}(k) - \mathbf{y})$, $\nabla_{\mathbf{v}}L(k) = \Phi(k)^T(\hat{\mathbf{y}}(k) - \mathbf{y})$ and $\lambda_{\min}(\mathbf{G}(k)) \geq \lambda_0/2$, we get

$$\begin{aligned} \langle \mathbf{g} - \hat{\mathbf{y}}(k), \hat{\mathbf{y}}(k) - \mathbf{y} \rangle &= -\alpha(\hat{\mathbf{y}}(k) - \mathbf{y})^T [\mathbf{Z}(k)\mathbf{Z}(k)^T + \Phi(k)\Phi(k)^T] (\hat{\mathbf{y}}(k) - \mathbf{y}) \\ &\leq -\frac{\alpha\lambda_0}{2} \|\hat{\mathbf{y}}(k) - \mathbf{y}\|^2. \end{aligned} \quad (43)$$

By combining the inequalities (41), 42, 43, we obtain

$$\begin{aligned} \|\hat{\mathbf{y}}(k+1) - \mathbf{y}\|^2 &\leq (1 - \alpha [\lambda_0 - C_0 c^2 n - \alpha C_0^2 c^4 n^2]) \beta^{2k} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \\ &\leq \left(1 - \frac{\alpha\lambda_0}{2}\right) \cdot \beta^{2k} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2 \\ &= \left(1 - \frac{\alpha\lambda_0}{2}\right)^{k+1} \|\hat{\mathbf{y}}(0) - \mathbf{y}\|^2, \end{aligned}$$

where the second inequality is by $\alpha = \mathcal{O}\left(\frac{\lambda_0}{n^2}\right)$. This proves the result (vi) and complete the proof. \square