

DICHOTOMY OF EARLY AND LATE PHASE IMPLICIT BIASES CAN PROVABLY INDUCE GROKKING

Kaifeng Lyu*

Princeton University
klyu@cs.princeton.edu

Jikai Jin*

Stanford University
jkjin@stanford.edu

Zhiyuan Li

Toyota Technological Institute at Chicago
zhiyuanli@ttic.edu

Simon S. Du

University of Washington
ssdu@cs.washington.edu

Jason D. Lee

Princeton University
jasonlee@princeton.edu

Wei Hu

University of Michigan
vvh@umich.edu

ABSTRACT

Recent work by Power et al. (2022) highlighted a surprising “grokking” phenomenon in learning arithmetic tasks: a neural net first “memorizes” the training set, resulting in perfect training accuracy but near-random test accuracy, and after training for sufficiently longer, it suddenly transitions to perfect test accuracy. This paper studies the grokking phenomenon in theoretical setups and shows that it can be induced by a dichotomy of early and late phase implicit biases. Specifically, when training homogeneous neural nets with large initialization and small weight decay on both classification and regression tasks, we prove that the training process gets trapped at a solution corresponding to a kernel predictor for a long time, and then a very sharp transition to min-norm/max-margin predictors occurs, leading to a dramatic change in test accuracy.

1 INTRODUCTION

The generalization behavior of modern over-parameterized neural nets has been puzzling: these nets have the capacity to overfit the training set, and yet they frequently exhibit a small gap between training and test performance when trained by popular gradient-based optimizers. A common view now is that the network architectures and training pipelines can automatically induce regularization effects to avoid or mitigate overfitting throughout the training trajectory.

Recently, Power et al. (2022) discovered an even more perplexing generalization phenomenon called *grokking*: when training a neural net to learn modular arithmetic operations, it first “memorizes” the training set with zero training error and near-random test error, and then training for much longer leads to a sharp transition from no generalization to perfect generalization. See Section 2 for our reproduction of this phenomenon. Beyond modular arithmetic, grokking has been reported in learning group operations (Chughtai et al., 2023), learning sparse parity (Barak et al., 2022; Bhattamishra et al., 2023), learning greatest common divisor (Charton, 2024), and image classification (Liu et al., 2023; Radhakrishnan et al., 2022).

Different viewpoints on the mechanism of grokking have been proposed, including the slingshot mechanism (cyclic phase transitions) (Thilak et al., 2022), random walk among minimizers (Milledge, 2022), slow formulation of good representations (Liu et al., 2022), the scale of initialization (Liu et al., 2023), and the simplicity of the generalizable solution (Nanda et al., 2023; Varma et al., 2023). However, existing studies failed to address two crucial aspects for gaining a comprehensive understanding of grokking:

1. No prior work has rigorously proved grokking in a neural network setting.
2. No prior work has provided a quantitative explanation as to why the transition from memorization to generalization is often sharp, instead of gradual.

*Equal Contribution

¹Code is available at <https://github.com/vfleaking/grokking-dichotomy>

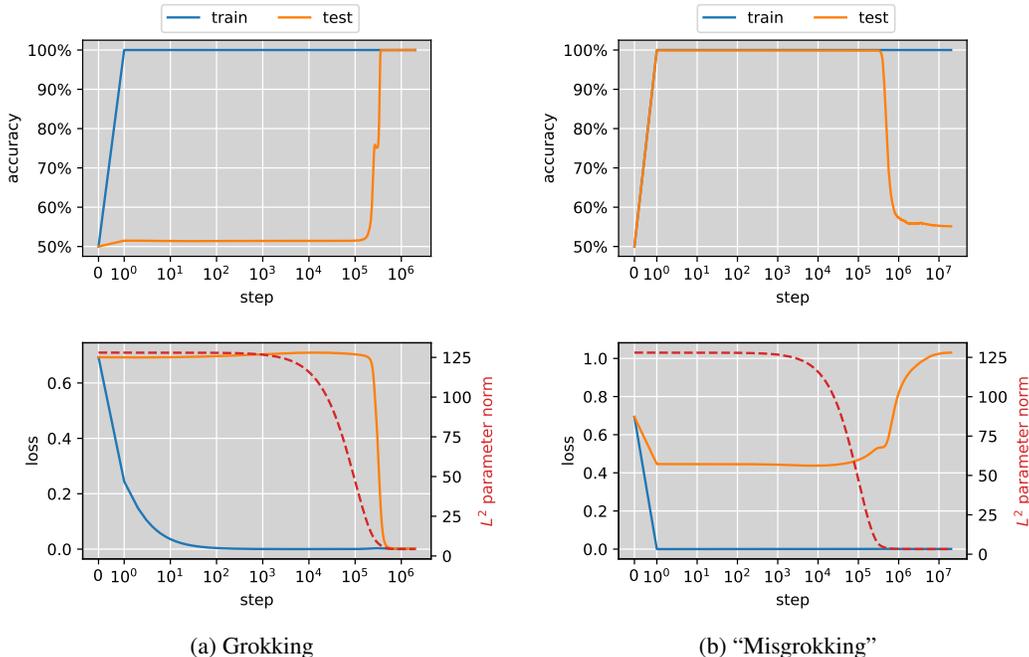


Figure 1: Training two-layer diagonal linear nets for linear classification can exhibit (a) grokking, for a dataset that can be linearly separated by a 3-sparse weight vector, or (b) “misgrokking”, for a dataset that can be linearly separated by a large L^2 -margin. See Section 3.2.3 for details.

Our Contributions. In this work, we address these limitations by identifying simple yet insightful theoretical setups where grokking with sharp transition can be rigorously proved and its mechanism can be intuitively understood. Our main intuition is that optimizers may implicitly induce different biases in early and late phases; grokking happens if the early phase bias implies an overfitting solution and late phase bias implies a generalizable solution.

More specifically, we focus on neural nets with large initialization and small weight decay. This is inspired by a recent work (Liu et al., 2023) showing that training with these two tricks can induce grokking on many tasks, even beyond learning modular arithmetic. Our theoretical analysis attributes this to the following *dichotomy* of early and late phase implicit biases: the large initialization induces a very strong early phase implicit bias towards kernel predictors, but over time, it decays and competes with a late phase implicit bias towards min-norm/max-margin predictors induced by the small weight decay, resulting in a transition that turns out to be provably sharp in between.

This implicit bias result holds for homogeneous neural nets, a broad class of neural nets that include commonly used MLPs and CNNs with homogeneous activation. We further exemplify this dichotomy in sparse linear classification with diagonal linear nets and low-rank matrix completion with over-parameterized models, where we provide theory and experiments showing that the kernel predictor does not generalize well but the min-norm/max-margin predictors do, hence grokking happens (see Figures 1a and 3).

Based on these theoretical insights, we are able to construct examples where the early phase implicit bias leads to a generalizable solution, and the late phase bias leads to overfitting. This gives rise to a new phenomenon which we call “*misgrokking*”: a neural net first achieves perfect training and test accuracies, but training for a longer time leads to a sudden big drop in test accuracy. See Figure 1b.

Our proof of the kernel regime in the early phase is related to the Neural Tangent Kernel (NTK) regime studied in the literature (Jacot et al., 2018; Chizat et al., 2019). However, our result is qualitatively different from existing NTK analyses in that in our setting, the weight norm changes significantly (due to weight decay), while in the usual NTK regime, the weight changes only by a small amount, and so does its norm. Our proof relies on a careful analysis of the norm and direction of the weight, which enables us to obtain a tight bound on the time spent in the kernel regime. This novel result of a kernel regime in homogeneous neural nets under weight decay may be of independent interest. We further analyze the late phase using a gradient convergence argument, and it turns out that a slightly longer time suffices for convergence to the max-margin/min-norm predictor, thus exhibiting a sharp transition between the two phases.

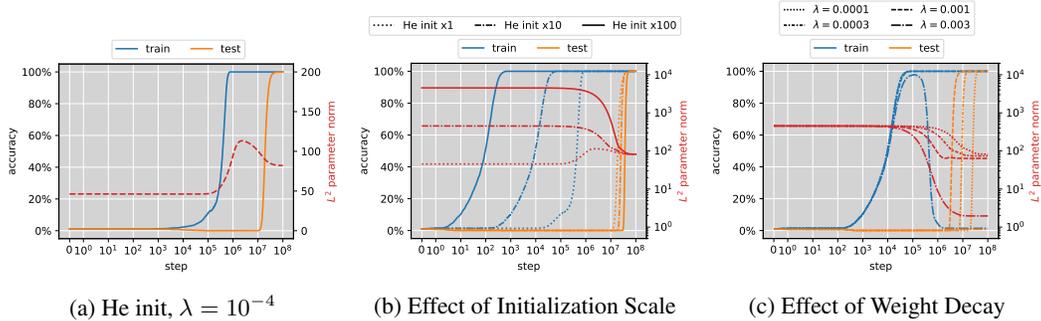


Figure 2: Training two-layer ReLU nets for modular addition exhibits grokking with He initialization and weight decay. Enlarging the initialization scale or reducing the weight decay delays the sharp transition in test accuracy.

Concretely, our contributions can be summarized as follows:

1. For training homogeneous neural nets with large initialization and small weight decay on classification and regression tasks, we prove a very sharp transition from optimal solutions of kernel SVM/regression to KKT solutions of global max-margin/min-norm problems associated with the neural net.
2. For classification and regression tasks respectively, we provide concrete examples with diagonal linear nets and overparameterized matrix completion, showing either grokking or misgrokking.

2 MOTIVATING EXPERIMENT: GROKING IN MODULAR ADDITION

In this section, we provide experiments to show that the initialization scale and weight decay are important factors in the grokking phenomenon, in the sense that enlarging the initialization scale or reducing the weight decay delays the sharp transition in test accuracy.

Grokking in Modular Addition. Following many previous works (Nanda et al., 2023; Gromov, 2023), we focus on *modular addition*, which belongs to the modular arithmetic tasks where the grokking phenomenon was first observed (Power et al., 2022). The task is to learn the addition operation over \mathbb{Z}_p : randomly split $\{(a, b, c) : a + b \equiv c \pmod{p}\}$ into training and test sets, and train a neural net on the training set to predict c given input pair (a, b) . The grokking phenomenon has been observed in learning this task under many training regimes, including training transformers and MLPs on cross-entropy and squared loss with various optimizers (full-batch GD, SGD, Adam, AdamW, etc.) (Power et al., 2022; Liu et al., 2022; Thilak et al., 2022; Gromov, 2023). For simplicity, we focus on a simple setting without involving too many confounding factors: training a two-layer ReLU net with full-batch GD. More specifically, we represent the input $(a, b) \in \mathbb{Z}_p \times \mathbb{Z}_p$ as the concatenation of the one-hot representations of a and b , resulting in an input vector $\mathbf{x} \in \mathbb{R}^{2p}$, and then the neural net processes the input as follows:

$$f(\boldsymbol{\theta}; \mathbf{x}) = \mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1), \quad (1)$$

where $\boldsymbol{\theta} = (\mathbf{W}_1, \mathbf{b}_1, \mathbf{W}_2)$ is the parameter vector consisting of $\mathbf{W}_1 \in \mathbb{R}^{h \times 2p}$, $\mathbf{b}_1 \in \mathbb{R}^h$, $\mathbf{W}_2 \in \mathbb{R}^{p \times h}$, h is the width, and $\text{ReLU}(v)$ stands for the ReLU activation that maps each entry v_i of a vector \mathbf{v} to $\max\{v_i, 0\}$. The neural net output $f(\boldsymbol{\theta}; \mathbf{x}) \in \mathbb{R}^p$ is treated as logits and is fed into the standard cross-entropy loss in training. We set the modulus p as 97, the training data size as 40% of the total number of data (p^2), the width h as 1024, the learning rate as 0.002 and weight decay as 10^{-4} . Figure 2a successfully reproduces the grokking phenomenon: the training accuracy reaches 100% in 10^6 steps while the test accuracy remains close to random guessing, but after training for one additional order of magnitude in the number of steps, the test accuracy suddenly jumps to 100%.

A natural question is how this grokking phenomenon depends on the training tricks in the pipeline. Liu et al. (2023) showed that, empirically, using a large initialization scale and a small but non-zero weight decay can induce grokking on various tasks even beyond modular arithmetic, including image classification on MNIST and sentiment classification on IMDB. Indeed, our experiments on modular addition can confirm the importance of initialization scale and weight decay.

Effect of Initialization Scale. To study the effect of initialization, we scale up the standard He initialization (He et al., 2015) by a factor of $\alpha > 0$, and then run the same training procedure. However, the model at large initialization produces a very large random number for each training

sample, which induces training instability. To remove this confounding factor, we set the weights W_2 in the second layer to $\mathbf{0}$ so that the initial model output is always 0. Results in Figure 2b show that the sharp transition in test accuracy is delayed when α increases from 1 to 100.

Effect of Weight Decay. In Figure 2c, we run experiments with different weight decay values. We observe that increasing weight decay makes the transition happen earlier. However, when weight decay surpasses a small threshold, the regularization strength can become so strong that it even hinders the optimization of the training loss and causes the training accuracy to collapse eventually. This collapse due to weight decay is also consistent with the observation in Lewkowycz and Gur-Ari (2020). Conversely, as weight decay decreases, the grokking phenomenon becomes more significant because now it takes more time for the test accuracy to exhibit the sharp transition. But the weight decay has to be non-zero: we will discuss in Appendix B.1 that training without weight decay leads to perfect generalization in the end, but the transition in test accuracy is no longer sharp.

Explanation? Liu et al. (2023) built their intuition upon an empirical observation from Fort and Scherlis (2019): there is a narrow range of weight norm, called *Goldilocks zone*, where the generalization inside the zone is better than the outside. It is argued that the weight norm can take a longer time to enter this Goldilocks zone when training with a larger initial weight norm or smaller weight decay, hence causing the delayed generalization. However, Fort and Scherlis (2019); Liu et al. (2023) did not provide much explanation on why such a narrow Goldilocks zone could exist in the first place. When talking about MLPs and CNNs with ReLU activation, which are commonly used in practice, this becomes even more mysterious: the output functions $f(\theta; \mathbf{x})$ of these neural nets are homogeneous to their parameters, i.e., $f(c\theta; \mathbf{x}) = c^L f(\theta; \mathbf{x})$ for all $c > 0$ (Lyu and Li, 2020; Ji and Telgarsky, 2020a; Kunin et al., 2023). As a result, it is impossible to explain the effect of norm just from the expressive power, since all classifiers that can be represented with a certain norm can also be represented by all other norms. This motivates us to dive deep into the training dynamics in grokking via rigorous theoretical analysis.

3 GROKING WITH LARGE INITIALIZATION AND SMALL WEIGHT DECAY

In this section, we present our theory on homogeneous neural nets with large initialization and small weight decay, which attributes grokking to a dichotomy of early and late phase implicit biases.

3.1 THEORETICAL SETUP

We focus on training models for classification and regression. Let \mathcal{D}_X be the input distribution. For every input $\mathbf{x} \sim \mathcal{D}_X$, let $y^*(\mathbf{x})$ be the classification/regression target of \mathbf{x} . We parameterize the model with θ and use $f(\theta; \mathbf{x})$ to denote the model output on input \mathbf{x} . We train the model by optimizing an empirical loss on a dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ where the data points are drawn i.i.d. as $\mathbf{x}_i \sim \mathcal{D}_X, y_i = y^*(\mathbf{x}_i)$. For each $i \in [n]$, we write $f_i(\theta) := f(\theta; \mathbf{x}_i)$ for short. We assume the model is *homogeneous* with respect to its parameters θ , a common assumption in analyzing neural networks (Lyu and Li, 2020; Ji and Telgarsky, 2020a; Telgarsky, 2023; Woodworth et al., 2020).

Assumption 3.1 For all \mathbf{x} , $f(\theta; \mathbf{x})$ is L -homogeneous (i.e., $f_i(c\theta) = c^L f_i(\theta)$ for all $c > 0$) and \mathcal{C}^2 -smooth with respect to θ .

The two-layer net in our motivating experiment and other MLPs and CNNs with ReLU activation are indeed L -homogeneous. However, we also need to assume the \mathcal{C}^2 -smoothness to ease the definition and analysis of gradient flow, which excludes the use of any non-smooth activation. This can be potentially addressed by defining and analyzing gradient flow via Clarke’s subdifferential (Clarke, 1975; 2001; Lyu and Li, 2020; Ji and Telgarsky, 2020a), but we choose to make this assumption for simplicity. Note that our concrete examples of grokking on sparse linear classification and matrix completion indeed satisfy this assumption.

As motivated in Section 2, we consider training with large initialization and small weight decay. Mathematically, we start training from $\alpha\theta_{\text{init}}$, where $\theta_{\text{init}} \in \mathbb{R}^d$ is a fixed vector and α is a large factor controlling the initialization scale. We also assume a small but non-zero weight decay $\lambda > 0$. To study the asymptotics more conveniently, we regard λ as a function of α with order $\lambda(\alpha) = \Theta(\alpha^{-p})$ for some positive $p = \Theta(1)$.

To avoid loss explosion at initialization, we deterministically set or randomly sample $\bar{\theta}_{\text{init}}$ in a way that the initial output of the model is zero. This can be done by setting the last layer weights to zero

(same as our experiments in Section 2), using the symmetrized initialization in Chizat et al. (2019), or using the “difference trick” in Hu et al. (2020). This zero-output initialization is also used in many previous studies of implicit bias (Chizat et al., 2019; Woodworth et al., 2020; Moroshko et al., 2020) for the sake of mathematical simplicity. We note that our analysis should be easily extended to random initialization with small outputs, just as in Chizat et al. (2019); Arora et al. (2019b).

Assumption 3.2 $f(\bar{\theta}_{\text{init}}; \mathbf{x}) = 0$ for all \mathbf{x} .

We define the vanilla training loss as the average over the loss values of each individual sample:

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(f_i(\theta); y_i). \quad (2)$$

With weight decay λ , it becomes the following regularized training loss:

$$\mathcal{L}_\lambda(\theta) = \mathcal{L}(\theta) + \frac{\lambda}{2} \|\theta\|_2^2. \quad (3)$$

For simplicity, in this paper, we consider minimizing the loss via gradient flow $\frac{d\theta}{dt} = -\nabla \mathcal{L}_\lambda(\theta)$, which is the continuous counterpart of gradient descent and stochastic gradient descent when the learning rate goes to 0. We use $\theta(t; \theta_0)$ to denote the parameter at time t by running gradient flow from θ_0 , and we write $\theta(t)$ when the initialization θ_0 is clear from the context.

3.2 GROKING IN CLASSIFICATION

In this subsection, we study the grokking phenomenon on classification tasks. For simplicity, we restrict to binary classification problems, where $y^*(\mathbf{x}) \in \{-1, +1\}$. While in practice it is standard to use the logistic loss $\ell(\hat{y}, y) = \log(1 + e^{-y\hat{y}})$ (a.k.a. binary cross-entropy loss), we follow the implicit bias literature to analyze the exponential loss $\ell(\hat{y}; y) = e^{-y\hat{y}}$ as a simple surrogate, which has the same tail behavior as the logistic loss when $y\hat{y} \rightarrow +\infty$ and thus usually leads to the same implicit bias (Soudry et al., 2018; Nacson et al., 2019b; Lyu and Li, 2020; Chizat and Bach, 2020).

By carefully analyzing the training dynamics, we rigorously prove that there is a sharp transition around $\frac{1}{\lambda} \log \alpha$. Before this point, gradient flow fits the training data perfectly while maintaining the parameter direction within a local region near the initial direction, which causes the classifier to behave like a kernel classifier based on Neural Tangent Kernel (NTK) (Jacot et al., 2018; Arora et al., 2019b;c). After the transition, however, the gradient flow escapes the local region and makes efforts to maximize the margin. Following the nomenclature in (Moroshko et al., 2020), we call the first regime the *kernel regime* and the second regime the *rich regime*. The key difference to the existing works analyzing the kernel and rich regimes (Moroshko et al., 2020; Telgarsky, 2023) is that the transition in our case is provably sharp, which is a crucial ingredient for the grokking phenomenon.

3.2.1 KERNEL REGIME

First, we show that gradient flow gets stuck in the kernel regime over the initial period of $\frac{1-c}{\lambda} \log \alpha$, where the model behaves as if it were optimizing over a linearized model. More specifically, for all \mathbf{x} , we define $\nabla f(\bar{\theta}_{\text{init}}; \mathbf{x})$ as the NTK feature of \mathbf{x} . As we are considering over-parameterized models, where the dimension of θ is larger than the number of data n , it is natural to assume that the NTK features of the training data are linearly separable (Ji and Telgarsky, 2020b; Telgarsky, 2023):

Assumption 3.3 *There exists \mathbf{h} such that $y_i \langle \nabla f_i(\bar{\theta}_{\text{init}}), \mathbf{h} \rangle > 0$ for all $i \in [n]$.*

Let $\mathbf{h}_{\text{ntk}}^*$ be the unit vector so that a linear classifier with weight $\mathbf{h}_{\text{ntk}}^*$ can attain the max L^2 -margin on the NTK features of the training data, $\{(\nabla f_i(\bar{\theta}_{\text{init}}), y_i)\}_{i=1}^n$, and γ_{ntk} be the corresponding L^2 -margin. That is, $\mathbf{h}_{\text{ntk}}^*$ is the unique unit vector that points to the direction of the unique optimal solution to the following constrained optimization problem:

$$\min \frac{1}{2} \|\mathbf{h}\|_2^2 \quad \text{s.t.} \quad y_i \langle \nabla f_i(\bar{\theta}_{\text{init}}), \mathbf{h} \rangle \geq 1, \quad \forall i \in [n], \quad (\text{K1})$$

and $\gamma_{\text{ntk}} := \min_{i \in [n]} \{y_i \langle \nabla f_i(\bar{\theta}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle\}$.

The following theorem states that for any $c \in (0, 1)$, the solution found by gradient flow at time $\frac{1-c}{\lambda} \log \alpha$ represents the same classifier as the max L^2 -margin linear classifier on the NTK features:

Theorem 3.4 For any all constants $c \in (0, 1)$, letting $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$, it holds that

$$\forall \mathbf{x} \in \mathbb{R}^d : \quad \lim_{\alpha \rightarrow +\infty} \frac{1}{Z(\alpha)} f(\boldsymbol{\theta}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}}); \mathbf{x}) = \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h}_{\text{ntk}}^* \rangle,$$

where $Z(\alpha) := \frac{1}{\gamma_{\text{ntk}}} \log \frac{\alpha^c}{\lambda}$ is a normalizing factor.

Key Proof Insight. The standard NTK-based dynamical analysis requires $\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) \approx \alpha \bar{\boldsymbol{\theta}}_{\text{init}}$, but the weight decay in our case can significantly change the norm, hence violating this condition. The key ingredient of our proof is a much more careful dynamical analysis showing that $\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) \approx \alpha e^{-\lambda t} \bar{\boldsymbol{\theta}}_{\text{init}}$, i.e., the norm of $\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})$ decays but the direction remains close to $\bar{\boldsymbol{\theta}}_{\text{init}}$. Then by L -homogeneity and Taylor expansion, we have

$$\begin{aligned} f(\boldsymbol{\theta}; \mathbf{x}) &= (\alpha e^{-\lambda t})^L f\left(\frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta}; \mathbf{x}\right) \approx (\alpha e^{-\lambda t})^L \left(f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}) + \left\langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}} \right\rangle \right) \\ &= \alpha^L e^{-L\lambda t} \left\langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}} \right\rangle. \end{aligned}$$

The large scaling factor $\alpha^L e^{-L\lambda t}$ above enables fitting the dataset even when the change in direction $\frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}}$ is very small. Indeed, we show that $\frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}} \approx \frac{1}{\gamma_{\text{ntk}}} \alpha^L e^{-L\lambda t} \log \frac{\alpha^c}{\lambda} \mathbf{h}_{\text{ntk}}^*$ at time $\frac{1-c}{\lambda} \log \alpha$ by closely tracking the dynamics, hence completing the proof. See [Appendix C.2](#) for details. When $\alpha^L e^{-L\lambda t}$ is no longer a large scaling factor, namely $t = \frac{1}{\lambda}(\log \alpha + \omega(1))$, this analysis breaks and the dynamics enter the rich regime.

3.2.2 RICH REGIME

Next, we show that continuing the gradient flow for a slightly long time to time $\frac{1+c}{\lambda} \log \alpha$, it is able to escape the kernel regime. More specifically, consider the following constrained optimization problem that aims at maximizing the margin of the predictor $f(\boldsymbol{\theta}; \mathbf{x})$ on the training data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$:

$$\min \frac{1}{2} \|\boldsymbol{\theta}\|_2^2 \quad \text{s.t.} \quad y_i f_i(\boldsymbol{\theta}) \geq 1, \quad \forall i \in [n]. \quad (\text{R1})$$

Then we have the following directional convergence result:

Theorem 3.5 For all constants $c > 0$ and for every sequence $\{\alpha_k\}_{k \geq 1}$ with $\alpha_k \rightarrow +\infty$, letting $T_c^+(\alpha) := \frac{1+c}{\lambda} \log \alpha$, there exists a time sequence $\{t_k\}_{k \geq 1}$ such that $\frac{1}{\lambda} \log \alpha_k \leq t_k \leq T_c^+(\alpha_k)$ and every limit point of $\left\{ \frac{\boldsymbol{\theta}(t_k; \alpha_k \bar{\boldsymbol{\theta}}_{\text{init}})}{\|\boldsymbol{\theta}(t_k; \alpha_k \bar{\boldsymbol{\theta}}_{\text{init}})\|_2} : k \geq 1 \right\}$ is along the direction of a KKT point of (R1).

For the problem (R1), the KKT condition is a first-order necessary condition for global optimality, but it is not sufficient in general since $f_i(\boldsymbol{\theta})$ can be highly non-convex. Nonetheless, since gradient flow can easily get trapped at spurious minima with only first-order information, the KKT condition is widely adopted as a surrogate for global optimality in theoretical analysis ([Lyu and Li, 2020](#); [Wang et al., 2021](#); [Kunin et al., 2023](#)).

Key Proof Insight. The key is to use the loss convergence and norm decay bounds from the kernel regime as a starting point to obtain a small upper bound for the gradient $\nabla \mathcal{L}_\lambda(\boldsymbol{\theta})$. Then we connect the gradient upper bounds with KKT conditions. See [Appendix C.3](#) for the proof.

3.2.3 EXAMPLE: LINEAR CLASSIFICATION WITH DIAGONAL LINEAR NETS

Now, we exemplify how our implicit bias results imply a sharp transition in test accuracy in a concrete setting: training diagonal linear nets for linear classification. Let \mathcal{D}_X be a distribution over \mathbb{R}^d and $y^*(\mathbf{x}) = \text{sign}(\langle \mathbf{w}^*, \mathbf{x} \rangle) \in \{\pm 1\}$ be the ground-truth target for binary classification, where $\mathbf{w}^* \in \mathbb{R}^d$ is an unknown vector. Following [Moroshko et al. \(2020\)](#), we consider training a so-called two-layer ‘‘diagonal’’ linear net to learn y^* . On input $\mathbf{x} \in \mathbb{R}^d$, the model outputs the following function $f(\boldsymbol{\theta}; \mathbf{x})$, where $\boldsymbol{\theta} := (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^d \times \mathbb{R}^d \simeq \mathbb{R}^{2d}$ is the model parameter.

$$f(\boldsymbol{\theta}; \mathbf{x}) = \sum_{k=1}^d u_k^2 x_k - \sum_{k=1}^d v_k^2 x_k. \quad (4)$$

This model is considered a diagonal net because it can be seen as a sparsely-connected two-layer net with $2d$ hidden neurons, each of which only inputs and outputs a single scalar. More specifically, the input \mathbf{x} is first expanded to an $2d$ -dimensional vector $(\mathbf{x}, -\mathbf{x})$. Then the k -th hidden neuron

in the first half inputs x_k and outputs $u_k x_k$, and the k -th hidden neuron in the second half inputs $-x_k$ and outputs $v_k x_k$. The output neuron in the second layer shares the same weights as the hidden neurons and takes a weighted sum over the hidden neuron outputs according to their weights: $f(\boldsymbol{\theta}; \mathbf{x}) = \sum_{k=1}^d u_k \cdot (u_k \cdot x_k) + \sum_{k=1}^d v_k \cdot (v_k \cdot (-x_k))$. Note that the model output is always linear with the input. For convenience, we can write $f(\boldsymbol{\theta}; \mathbf{x}) = \langle \mathbf{w}(\boldsymbol{\theta}), \mathbf{x} \rangle$ with effective weight $\mathbf{w}(\boldsymbol{\theta}) := \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}$.

In this setup, we can understand the early and late phase implicit biases very concretely. We start training from a large initialization: $\boldsymbol{\theta}(0) = \alpha \boldsymbol{\theta}_{\text{init}}$, where $\alpha > 0$ controls the initialization scale and $\boldsymbol{\theta}_{\text{init}} = (\mathbf{1}, \mathbf{1})$. That is, $u_k = v_k = \alpha$ for all $1 \leq k \leq d$ at $t = 0$. As noted in Moroshko et al. (2020), the kernel feature for each data point is $\nabla f(\boldsymbol{\theta}_{\text{init}}; \mathbf{x}) = (2\mathbf{x}, -2\mathbf{x})$, so the max L^2 -margin kernel classifier is the max L^2 -margin linear classifier. This leads to the following corollary:

Corollary 3.6 (Diagonal linear nets, kernel regime) *Let $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$. For all constants $c \in (0, 1)$, as $\alpha \rightarrow +\infty$, the normalized effective weight vector $\frac{\mathbf{w}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha \boldsymbol{\theta}_{\text{init}}))}{\|\mathbf{w}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha \boldsymbol{\theta}_{\text{init}}))\|_2}$ converges to the max L^2 -margin direction, namely $\mathbf{w}_2^* := \arg \max_{\mathbf{w} \in \mathcal{S}^{d-1}} \{\min_{i \in [n]} y_i \langle \mathbf{w}, \mathbf{x}_i \rangle\}$.*

In contrast, if we train slightly longer from $T_c^-(\alpha)$ to $T_c^+(\alpha)$, we can obtain a KKT point of (R1).

Corollary 3.7 (Diagonal linear nets, rich regime) *Let $T_c^+(\alpha) := \frac{1+c}{\lambda} \log \alpha$. For all constants $c \in (0, 1)$, and for every sequence $\{\alpha_k\}_{k \geq 1}$ with $\alpha_k \rightarrow +\infty$, there exists a time sequence $\{t_k\}_{k \geq 1}$ such that $\frac{1}{\lambda} \log \alpha_k \leq t_k \leq T_c^+(\alpha_k)$ and every limit point of $\left\{ \frac{\boldsymbol{\theta}(t_k; \alpha_k \boldsymbol{\theta}_{\text{init}})}{\|\boldsymbol{\theta}(t_k; \alpha_k \boldsymbol{\theta}_{\text{init}})\|_2} : k \geq 1 \right\}$ is along the direction of a KKT point of the problem $\min \|\mathbf{u}\|_2^2 + \|\mathbf{v}\|_2^2$ s.t. $y_i \langle \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}, \mathbf{x}_i \rangle \geq 1$.*

As noted in these two works, in the diagonal linear case, optimizing (R1) to the global optimum is equivalent to finding the max L^1 -margin linear classifier: $\min \|\mathbf{w}\|_1$ s.t. $y_i \langle \mathbf{w}, \mathbf{x}_i \rangle \geq 1, \forall i \in [n]$, or equivalently, $\mathbf{w}_1^* := \arg \max_{\mathbf{w} \in \mathcal{S}^{d-1}} \{\min_{i \in [n]} y_i \langle \mathbf{w}, \mathbf{x}_i \rangle\}$. Therefore, our theory suggests a sharp transition at time $\frac{1}{\lambda} \log \alpha$ from max L^2 -margin linear classifier to max L^1 -margin linear classifier.

While this corollary only shows KKT conditions rather than global optimality, we believe that one may be able to obtain the global optimality using insights from existing analysis of training diagonal linear nets without weight decay (Gunasekar et al., 2018b; Moroshko et al., 2020).

Empirical Validation: Grokking. As it is well known that maximizing the L^1 -margin can better encourage sparsity than maximizing the L^2 -margin, our theory predicts that the grokking phenomenon can be observed in training diagonal linear nets for k -sparse linear classification. To verify this, we specifically consider the following task: sample n data points uniformly from $\{\pm 1\}^d$ for a very large d , and let the ground truth be a linear classifier with a k -sparse weight, where the non-zero coordinates are sampled uniformly from $\{\pm 1\}^k$. Applying standard generalization bounds based on Rademacher complexity can show that the max L^2 -margin linear classifier needs $\mathcal{O}(kd)$ to generalize, while the max L^1 -margin linear classifier only needs $\mathcal{O}(k^2 \log d)$. See Appendix E.2 for the proof. This suggests that the grokking phenomenon can be observed in this setting. In Figure 1a, we run experiments with $n = 256, d = 10^5, k = 3$, large initialization with initial parameter norm 128, and small weight decay $\lambda = 0.001$. As expected, we observe the grokking phenomenon: in the early phase, the net fits the training set very quickly but fails to generalize; after a sharp transition in the late phase, the test accuracy becomes 100%.

Empirical Validation: Misgrokking. From the above example, we can see that grokking can be induced by the mismatch between the early-phase implicit bias and data and a good match between the late-phase implicit bias and data. However, this match and mismatch can largely depend on data. Conversely, if we consider the case where the labels in the linear classification problem are generated by a linear classifier with a large L^2 -margin, the early phase implicit bias can make the neural net generalize easily, but the late phase implicit bias can destroy this good generalization since the ground-truth weight vector may not have a large L^1 -margin. To justify this, we first sample a unit-norm vector \mathbf{w}^* , then sample a dataset from the distribution $(\mathbf{x}, y) \sim (\mathbf{z} + \frac{\gamma}{2} \text{sign}(\langle \mathbf{z}, \mathbf{w}^* \rangle) \mathbf{w}^*, \text{sign}(\langle \mathbf{z}, \mathbf{w}^* \rangle))$, where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ (i.e., a Gaussian distribution that is separated by a margin in the middle). We take $\gamma = 25$ and sample $n = 32$ points. Indeed, we observe a phenomenon which we call “misgrokking”: the neural net first fits the training set and achieves 100% test accuracy, and then after training for sufficiently longer, the test accuracy drops to nearly 50%. See Figure 1b.

3.3 GROKING IN REGRESSION

In this subsection, we study the grokking phenomenon on regression tasks with squared loss $\ell(\hat{y}, y) = (\hat{y} - y)^2$. Paralleling the classification setting, we analyze the behavior of gradient flow in both kernel and rich regimes, and show that there is a sharp transition around $\frac{1}{\lambda} \log \alpha$.

3.3.1 KERNEL REGIME

Similar to the classification setting, we first show that gradient flow gets stuck in the kernel regime over the initial period of $\frac{1-c}{\lambda} \log \alpha$ time. For all \mathbf{x} , we define $\nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x})$ as the NTK feature of \mathbf{x} . For over-parameterized models, where the dimension of $\boldsymbol{\theta}$ is larger than the number of data n , we make the following natural assumption that is also widely used in the literature (Du et al., 2019b; Chizat et al., 2019; Arora et al., 2019b):

Assumption 3.8 *The NTK features of training samples $\{\nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x})\}_{i=1}^n$ are linearly independent.*

Now we let $\mathbf{h}_{\text{ntk}}^*$ be the vector with minimum norm such that the linear predictor $\mathbf{g} \mapsto \langle \mathbf{g}, \mathbf{h} \rangle$ perfectly fits $\{(\nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), y_i)\}_{i=1}^n$. That is, $\mathbf{h}_{\text{ntk}}^*$ is the solution to the following constrained optimization problem:

$$\min \frac{1}{2} \|\mathbf{h}\|_2^2 \quad \text{s.t.} \quad \langle \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h} \rangle = y_i, \quad \forall i \in [n]. \quad (\text{K2})$$

Then we have the following result that is analogous to [Theorem 3.4](#). See [Appendix C.2](#) for the proof.

Theorem 3.9 *For all constants $c \in (0, 1)$, letting $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$, it holds that*

$$\forall \mathbf{x} \in \mathbb{R}^d: \quad \lim_{\alpha \rightarrow +\infty} f(\boldsymbol{\theta}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}}); \mathbf{x}) = \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h}_{\text{ntk}}^* \rangle.$$

3.3.2 RICH REGIME

Similar to the classification setting, we then show that gradient flow is able to escape the kernel regime at time $\frac{1+c}{\lambda} \log \alpha$. Specifically, consider the following constrained optimization problem that searches for the parameter with the minimum norm that can perfectly fit the training data:

$$\min \frac{1}{2} \|\boldsymbol{\theta}\|_2^2 \quad \text{s.t.} \quad f_i(\boldsymbol{\theta}) = y_i, \quad \forall i \in [n]. \quad (\text{R2})$$

Then we have the following convergence result. See [Appendix D.2](#) for the proof.

Theorem 3.10 *For any constant $c > 0$, letting $T_c^+(\alpha) := \frac{1+c}{\lambda} \log \alpha$, for any sequence of $\{\alpha_k\}_{k \geq 1}$ with $\alpha_k \rightarrow +\infty$, there exists a time sequence $\{t_k\}_{k \geq 1}$ satisfying $\frac{1}{\lambda} \log \alpha_k \leq t_k \leq T_c^+(\alpha_k)$, such that $\|\boldsymbol{\theta}(t_k; \alpha_k \bar{\boldsymbol{\theta}}_{\text{init}})\|_2$ are uniformly bounded and that every limit point of $\{\boldsymbol{\theta}(t_k; \alpha_k \bar{\boldsymbol{\theta}}_{\text{init}}) : k \geq 1\}$ is a KKT point of [\(R2\)](#).*

3.3.3 EXAMPLE: OVERPARAMETERIZED MATRIX COMPLETION

As an example, we consider the matrix completion problem of recovering a low-rank symmetric matrix based on partial observations of its entries. Specifically, we let $\mathbf{X}^* \in \mathbb{R}^{d \times d}$ be the ground-truth symmetric matrix with rank $r \ll d$ and we observe a (random) set of entries indexed by $\Omega \subseteq [d] \times [d]$. In this setting, the input distribution \mathcal{D}_X is the uniform distribution on the finite set $[d] \times [d]$. For every input $\mathbf{x} = (i, j) \sim \mathcal{D}_X$, $y^*(\mathbf{x}) = \langle \mathbf{P}_x, \mathbf{X}^* \rangle$ where $\mathbf{P}_x = \mathbf{e}_i \mathbf{e}_j^\top$.

Following previous works (Arora et al., 2019a; Razin and Cohen, 2020; Woodworth et al., 2020; Li et al., 2021), we solve the matrix completion problem using the matrix factorization approach, *i.e.*, we parameterize a matrix as $\mathbf{W} = \mathbf{U}\mathbf{U}^\top - \mathbf{V}\mathbf{V}^\top$, where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{d \times d}$, and optimize \mathbf{U}, \mathbf{V} so that \mathbf{W} matches with the ground truth on observed entries. This can be viewed as a neural net that is parameterized by $\boldsymbol{\theta} = (\text{vec}(\mathbf{U}), \text{vec}(\mathbf{V})) \in \mathbb{R}^{2d^2}$ and outputs $f(\boldsymbol{\theta}; \mathbf{x}) = \langle \mathbf{P}_x, \mathbf{W}(\boldsymbol{\theta}) \rangle$ given $\mathbf{x} = (i, j)$, where $\mathbf{W}(\boldsymbol{\theta}) := \mathbf{U}\mathbf{U}^\top - \mathbf{V}\mathbf{V}^\top$. It is easy to see that $f(\boldsymbol{\theta}; \mathbf{x})$ is 2-homogeneous, satisfying [Assumption 3.1](#). Given a dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, we consider using gradient flow to minimize the ℓ_2 -regularized square loss defined in [\(3\)](#) which can be equivalently written as

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n (f(\boldsymbol{\theta}; \mathbf{x}_i) - y_i)^2 + \frac{\lambda}{2} (\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2). \quad (5)$$

In the kernel regime, [Theorem 3.9](#) implies the following result for identity initialization.

Corollary 3.11 (Matrix completion, kernel regime) Let $\bar{\theta}_{\text{init}} := (\text{vec}(\mathbf{I}), \text{vec}(\mathbf{I}))$. For all constants $c \in (0, 1)$, letting $T_c^-(\alpha) = \frac{1-c}{\lambda} \log \alpha$, it holds for all $(i, j) \in [d] \times [d]$ that

$$\lim_{\alpha \rightarrow +\infty} [\mathbf{W}(\theta(T_c^-(\alpha); \alpha \bar{\theta}_{\text{init}}))]_{ij} = \begin{cases} \mathbf{X}_{ij}^* & \text{if } (i, j) \in \Omega \text{ or } (j, i) \in \Omega, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, in the kernel regime, while GD is able to fit the observed entries, it always fills the unobserved entries with 0, leading to a significant failure in recovering the matrix.

In the rich regime, [Theorem 3.10](#) implies that the solution transitions sharply near $\frac{1}{\lambda} \log \alpha$ to a KKT point of the norm minimization problem. If this KKT point is globally optimal, then it is known that $\mathbf{W}(\theta)$ is the minimum nuclear norm solution to the matrix completion problem: $\min \|\mathbf{W}\|_*$ s.t. $\mathbf{W} = \mathbf{W}^\top$, $\mathbf{W}_{ij} = \mathbf{X}_{ij}^*$, $\forall (i, j) \in \Omega$ ([Ding et al., 2022](#), Theorem 3.2). It is known that the minimum nuclear norm solution can recover the ground truth with very small errors when $\tilde{\mathcal{O}}(d \log^2 d)$ entries are observed ([Theorem F.5](#)).

Although it is not proved that this KKT point is globally optimal, the following theorem provides positive evidence by showing that gradient flow eventually converges global minima and recovers the ground truth, if the time is not limited around $\frac{1}{\lambda} \log \alpha$.

Theorem 3.12 (Matrix completion, rich regime) Suppose that $\text{rank}(\mathbf{X}^*) = r = \mathcal{O}(1)$ and $\mathbf{X}^* = \mathbf{V}_{\mathbf{X}^*} \Sigma_{\mathbf{X}^*} \mathbf{V}_{\mathbf{X}^*}^\top$ is a SVD of \mathbf{X}^* , where each row of $\mathbf{V}_{\mathbf{X}^*}$ has ℓ_∞ -norm bounded by $\sqrt{\frac{r}{d}}$. If the number of observed entries satisfies $N \gtrsim \mu^4 d \log^2 d$, then for any $\sigma > 0$, the gradient flow trajectory $(\theta(t))_{t \geq 0}$ for (5) starting from random initialization $(\mathbf{U}(0))_{ij}, (\mathbf{V}(0))_{ij} \stackrel{i.i.d.}{\sim} \mathcal{N}(\alpha \mathbb{1}\{i=j\}, \sigma^2)$ converges to a global minimizer θ_∞ of \mathcal{L}_λ . Moreover, $\|\mathbf{W}(\theta_\infty) - \mathbf{X}^*\|_F \lesssim \sqrt{\lambda} \|\mathbf{X}^*\|_* \mu^2 \log d$ with probability $\geq 1 - d^{-3}$.

Remark 3.13 The small random perturbation to the identity initialization $(\alpha \mathbf{I}, \alpha \mathbf{I})$ in the above theorem is needed to guarantee that gradient flow does not get stuck at saddle points and thus converges to global minimizers almost surely. By the continuity of the gradient flow trajectory w.r.t. initialization, the conclusion of [Corollary 3.11](#) still holds if $\sigma = \sigma(\alpha)$ is sufficiently small.

Empirical Validation. We verify our theory on a simple algorithmic task: matrix completion for the multiplication table. Here, the multiplication table refers to a matrix \mathbf{X}^* where the (i, j) -entry is ij/d^2 . It is easy to see that this matrix is of rank-1 since $\mathbf{X}^* = \mathbf{u}\mathbf{u}^\top$ for $\mathbf{u} = (0, \frac{1}{d}, \frac{2}{d}, \dots, \frac{d-1}{d})$. This means the late phase bias can complete the matrix with few observed entries; however, the early phase bias fills every unobserved entry with 0, leading to a large test loss. This dichotomy of implicit biases implies grokking. To check this, we set $d = 97$ and randomly choose 5% of the entries as the training set. We run GD with initialization scale $\alpha = 10$, learning rate $\eta = 0.1$, weight decay 10^{-4} . [Figure 3](#) shows that grokking indeed happens: GD takes $\sim 10^1$ steps to minimize the training loss, but the sharp transition in test loss occurs only after $\sim 10^6$ steps.

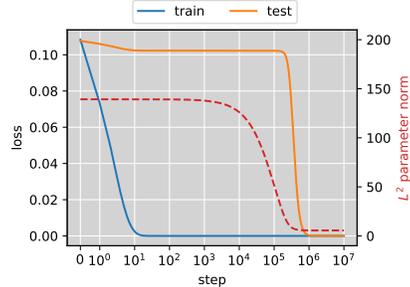


Figure 3: Matrix completion for the multiplication table exhibits grokking.

4 CONCLUSION

In this paper, we show that the grokking phenomenon provably occurs in several setups. Our results suggest that the sharp transition in test accuracy may stem from a dichotomy of implicit biases between the early and late training phases. While the early phase bias guides the optimizer to overfitting solutions, it is quickly corrected when the late phase bias takes effect.

Some limitations of our work are as follows. First, our work only studies the training dynamics with large initialization and small weight decay, but these may not be the only source of the dichotomy of the implicit biases. We further discuss in [Appendix B](#) that the late phase implicit bias can also be induced by implicit margin maximization and sharpness reduction, though the transition may not be as sharp as the weight decay case. Also, our work focuses on understanding the cause of grokking but does not study how to make neural nets generalize without so much delay in time. We leave it to future work to explore the other sources of grokking and practical methods to eliminate grokking.

ACKNOWLEDGMENTS

Kaifeng Lyu is partly supported by NSF and ONR. Simon S. Du is supported by supported by NSF IIS 2110170, NSF DMS 2134106, NSF CCF 2212261, NSF IIS 2143493, NSF CCF 2019844, NSF IIS 2229881. Wei Hu is supported by Google Research Scholar Program.

REFERENCES

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via overparameterization. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 242–252. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/allen-zhu19a.html>.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/c0c783b5fc0d7d808f1d14a6e9c8280d-Paper.pdf.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 322–332. PMLR, 09–15 Jun 2019b. URL <https://proceedings.mlr.press/v97/arora19a.html>.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019c. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/dbc4d84bfcfe2284ba11beffb853a8c4-Paper.pdf.
- Pranjal Awasthi, Natalie Frank, and Mehryar Mohri. On the Rademacher complexity of linear hypothesis sets. *arXiv preprint arXiv:2007.11045*, 2020.
- Boaz Barak, Benjamin Edelman, Surbhi Goel, Sham Kakade, Eran Malach, and Cyril Zhang. Hidden progress in deep learning: SGD learns parities near the computational limit. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 21750–21764. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/884ba65392170763b27c914087bde01-Paper-Conference.pdf.
- Satwik Bhattamishra, Arkil Patel, Varun Kanade, and Phil Blunsom. Simplicity bias in transformers and their ability to learn sparse boolean functions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5767–5791, 2023.
- Guy Blanc, Neha Gupta, Gregory Valiant, and Paul Valiant. Implicit regularization for deep neural networks driven by an Ornstein-Uhlenbeck like process. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 483–513. PMLR, 09–12 Jul 2020. URL <https://proceedings.mlr.press/v125/blanc20a.html>.
- Emmanuel J Candes and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/cf9dc5e4e194fc21f397b4cac9cc3ae9-Paper.pdf.

- Francois Charton. Learning the greatest common divisor: explaining transformer predictions. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=cmcD05NPKa>.
- Lenaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR, 2020.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/a614c557843b1df326cb29c57225459-Paper.pdf.
- Bilal Chughtai, Lawrence Chan, and Neel Nanda. A toy model of universality: Reverse engineering how networks learn group operations. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 6243–6267. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/chughtai23a.html>.
- Francis Clarke. Nonsmooth analysis in control theory: a survey. *European Journal of Control*, 7(2-3):145–159, 2001.
- Frank H Clarke. Generalized gradients and applications. *Transactions of the American Mathematical Society*, 205:247–262, 1975.
- Alex Damian, Tengyu Ma, and Jason D Lee. Label noise SGD provably prefers flat global minimizers. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 27449–27461. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/e6af401c28c1790eaf7d55c92ab6ab6-Paper.pdf.
- Xander Davies, Lauro Langosco, and David Krueger. Unifying grokking and double descent. In *NeurIPS ML Safety Workshop*, 2022. URL <https://openreview.net/forum?id=JqtHMZtqWm>.
- Lijun Ding, Dmitriy Drusvyatskiy, Maryam Fazel, and Zaid Harchaoui. Flat minima generalize for low-rank matrix recovery. *arXiv preprint arXiv:2203.03756*, 2022.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1675–1685. PMLR, 09–15 Jun 2019a. URL <https://proceedings.mlr.press/v97/dul9c.html>.
- Simon S. Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019b. URL <https://openreview.net/forum?id=S1eK3i09YQ>.
- Stanislav Fort and Adam Scherlis. The Goldilocks zone: Towards better understanding of neural network loss landscapes. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’19/IAAI’19/EAAI’19. AAAI Press, 2019. ISBN 978-1-57735-809-1. doi: 10.1609/aaai.v33i01.33013574. URL <https://doi.org/10.1609/aaai.v33i01.33013574>.
- Andrey Gromov. Grokking modular arithmetic. *arXiv preprint arXiv:2301.02679*, 2023.
- Xinran Gu, Kaifeng Lyu, Longbo Huang, and Sanjeev Arora. Why (and when) does local SGD generalize better than SGD? In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=svCcui6Drl>.

- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/58191d2a914c6dae66371c9dc91b41-Paper.pdf.
- Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1832–1841. PMLR, 10–15 Jul 2018a. URL <https://proceedings.mlr.press/v80/gunasekar18a.html>.
- Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018b. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/0e98aeeb54acf612b9eb4e48a269814c-Paper.pdf.
- Jeff Z HaoChen, Colin Wei, Jason Lee, and Tengyu Ma. Shape matters: Understanding the implicit bias of the noise covariance. In *Conference on Learning Theory*, pages 2315–2357. PMLR, 2021.
- 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 (ICCV)*, December 2015.
- Michael Y Hu, Angelica Chen, Naomi Saphra, and Kyunghyun Cho. Latent state models of training dynamics. *Transactions on Machine Learning Research*, 2023.
- Wei Hu, Zhiyuan Li, and Dingli Yu. Simple and effective regularization methods for training on noisily labeled data with generalization guarantee. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=Hke3gyHYwH>.
- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/5a4be1fa34e62bb8a6ec6b91d2462f5a-Paper.pdf.
- Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 17176–17186. Curran Associates, Inc., 2020a. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/c76e4b2fa54f8506719a5c0dc14c2eb9-Paper.pdf.
- Ziwei Ji and Matus Telgarsky. Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow relu networks. In *International Conference on Learning Representations*, 2020b. URL <https://openreview.net/forum?id=HygegyrYwH>.
- V. Koltchinskii and D. Panchenko. Empirical Margin Distributions and Bounding the Generalization Error of Combined Classifiers. *The Annals of Statistics*, 30(1):1 – 50, 2002. doi: 10.1214/aos/1015362183.
- Tanishq Kumar, Blake Bordelon, Samuel J. Gershman, and Cengiz Pehlevan. Grokking as the transition from lazy to rich training dynamics. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=vt5mnLVIVo>.
- Daniel Kunin, Atsushi Yamamura, Chao Ma, and Surya Ganguli. The asymmetric maximum margin bias of quasi-homogeneous neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=IM4xp7kGI5V>.

- Noam Itzhak Levi, Alon Beck, and Yohai Bar-Sinai. Grokking in linear estimators – a solvable model that groks without understanding. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=GH2LYb9XV0>.
- Aitor Lewkowycz and Guy Gur-Ari. On the training dynamics of deep networks with l2 regularization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4790–4799. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/32fcc8cfelfa4c77b5c58dafd36d1a98-Paper.pdf.
- Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=AH0s7Sm5H7R>.
- Zhiyuan Li, Tianhao Wang, and Sanjeev Arora. What happens after SGD reaches zero loss? –a mathematical framework. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=siCt4xZn5Ve>.
- Ziming Liu, Ouail Kitouni, Niklas S Nolte, Eric Michaud, Max Tegmark, and Mike Williams. Towards understanding grokking: An effective theory of representation learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 34651–34663. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/dfc310e81992d2e4cedc09ac47eff13e-Paper-Conference.pdf.
- Ziming Liu, Eric J Michaud, and Max Tegmark. Omnigrok: Grokking beyond algorithmic data. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=zDiHoIWa0ql>.
- Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SJeLIgBKPS>.
- Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 34689–34708. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/dfcd1c523512e557f4e75e8309049213-Paper-Conference.pdf.
- William Merrill, Nikolaos Tsilivis, and Aman Shukla. A tale of two circuits: Grokking as competition of sparse and dense subnetworks. In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2023.
- Beren Millidge. Grokking. 2022. URL <https://www.beren.io/2022-01-11-Grokking-g-Grokking/>.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT press, 2018.
- Edward Moroshko, Blake E Woodworth, Suriya Gunasekar, Jason D Lee, Nati Srebro, and Daniel Soudry. Implicit bias in deep linear classification: Initialization scale vs training accuracy. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 22182–22193. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/fc2022c89b61c76bbe978f1370660bf-Paper.pdf.
- Depen Morwani, Benjamin L. Edelman, Costin-Andrei Oncescu, Rosie Zhao, and Sham M. Kakade. Feature emergence via margin maximization: case studies in algebraic tasks. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=i9wDX850jR>.

- Mor Shpigel Nacson, Suriya Gunasekar, Jason Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4683–4692. PMLR, 09–15 Jun 2019a. URL <https://proceedings.mlr.press/v97/nacson19a.html>.
- Mor Shpigel Nacson, Jason Lee, Suriya Gunasekar, Pedro Henrique Pamplona Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 3420–3428. PMLR, 16–18 Apr 2019b. URL <https://proceedings.mlr.press/v89/nacson19b.html>.
- Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=9XF5bDPmdW>.
- Pascal Notsawo Jr, Hattie Zhou, Mohammad Pezeshki, Irina Rish, Guillaume Dumas, et al. Predicting grokking long before it happens: A look into the loss landscape of models which grok. *arXiv preprint arXiv:2306.13253*, 2023.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
- Adityanarayanan Radhakrishnan, Daniel Beaglehole, Parthe Pandit, and Mikhail Belkin. Mechanism of feature learning in deep fully connected networks and kernel machines that recursively learn features. *arXiv preprint arXiv:2212.13881*, 2022.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21174–21187. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f21e255f89e0f258accbe4e984eef486-Paper.pdf.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Matus Telgarsky. Feature selection and low test error in shallow low-rotation reLU networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=swEskiem99>.
- Vimal Thilak, Etai Littwin, Shuangfei Zhai, Omid Saremi, Roni Paiss, and Joshua Susskind. The slingshot mechanism: An empirical study of adaptive optimizers and the grokking phenomenon. *arXiv preprint arXiv:2206.04817*, 2022.
- Vikrant Varma, Rohin Shah, Zachary Kenton, János Kramár, and Ramana Kumar. Explaining grokking through circuit efficiency. *arXiv preprint arXiv:2309.02390*, 2023.
- Bohan Wang, Qi Meng, Wei Chen, and Tie-Yan Liu. The implicit bias for adaptive optimization algorithms on homogeneous neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10849–10858. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/wang21q.html>.
- G Alistair Watson. Characterization of the subdifferential of some matrix norms. *Linear algebra and its applications*, 170(0):33–45, 1992.
- Kaiyue Wen, Tengyu Ma, and Zhiyuan Li. How sharpness-aware minimization minimizes sharpness? In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=5spDgWmpY6x>.

Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, pages 3635–3673. PMLR, 2020.

Zhiwei Xu, Yutong Wang, Spencer Frei, Gal Vardi, and Wei Hu. Benign overfitting and grokking in reLU networks for XOR cluster data. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=BxHgpC6FNv>.

Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes over-parameterized deep ReLU networks. *Machine learning*, 109:467–492, 2020.

Bojan Žunkovič and Enej Ilievski. Grokking phase transitions in learning local rules with gradient descent. *arXiv preprint arXiv:2210.15435*, 2022.

A ADDITIONAL RELATED WORKS

Identifying Mechanisms that Cause Grokking. Some works attempted to identify mechanisms that cause grokking through experiments. [Davies et al. \(2022\)](#) hypothesized that both grokking and double descent are the result of the existence of two patterns: one pattern is faster to learn but generalizes poorly, and the other pattern is slower to learn but generalizes well. However, it is unclear how these notions can be rigorously defined for general training dynamics. Similar to the work of [Liu et al. \(2023\)](#) that we discussed in [Section 2](#), [Varma et al. \(2023\)](#) hypothesized that grokking happens when the generalizable solutions have a smaller parameter norm than the overfitting solutions, but the former takes a longer time to learn. They also demonstrated that violating a part of this condition indeed changes the generalization behavior from grokking to “ungrokking” and “semi-grokking”, two interesting phenomena they named in the paper. For training with adaptive gradient methods, [Thilak et al. \(2022\)](#) identified the Slingshot mechanism, which refers to a cyclic behavior of the last-layer parameter norm that empirically causes grokking. [Notsawo Jr et al. \(2023\)](#) also discussed a similar relationship between the oscillations in training loss and grokking. For the task of learning sparse parity, [Merrill et al. \(2023\)](#) empirically observed that the number of active neurons significantly decrease as the test accuracy starts to improve. While many explanations of grokking may make sense intuitively, none of these works provides rigorous justification their claims with mathematical analysis for neural net training, while our work is grounded by theoretical analyses of implicit biases in the kernel and rich regimes.

Progress Measures of Grokking. Efforts have been made to find progress measures that are improving smoothly and are predictive of the time to perfect generalization. [Nanda et al. \(2023\)](#); [Chughtai et al. \(2023\)](#); [Gromov \(2023\)](#) focused on the tasks of learning modular arithmetic or more general group operations, and attempted to reverse engineer the weights of neural nets after grokking. It was found that the weights can exhibit special structures to make the neural net internally compute trigonometric functions or other representations of group elements. These works also made efforts to define progress measures for grokking based on this understanding of the final weights. In a related study, [Morwani et al. \(2024\)](#) derived analytical formulas for the weights of two-layer neural nets, assuming the margin is ultimately maximized. [Hu et al. \(2023\)](#) computed a variety of metrics throughout training and fit a Hidden Markov Model (HMM) over them to study the phase transition in grokking. A common issue with these works is that why progress measures themselves can make progress is still not well understood and requires new insights into the training dynamics, which is the focus of our work.

Dynamical Analysis of Grokking. A few recent works were devoted to theoretically analyze the training dynamics in the grokking phenomenon. [Liu et al. \(2022\)](#) attributed the delayed generalization in grokking to the slow learning speed of the input embeddings. They especially analyzed the dynamics of the input embeddings in a related optimization problem, and showed that the eigenvalues of the coefficient matrix in the ODE can be used to predict the convergence. However, it is unclear how this relates to the training dynamics of the original problem and why the neural net overfits the dataset before the predicted convergence time. [Žunković and Ilievski \(2022\)](#) analyzed a simple linear classification problem, where the loss is the squared loss with explicit ℓ_1 and ℓ_2 regularization. Assuming specific input distributions, they proved sharp transition in test accuracy and derived an estimate of the grokking time, which depends on the initialization and regularization strength. [Levi et al. \(2024\)](#) leveraged random matrix theory to analyze grokking in a standard linear regression setting, where the test accuracy can provably exhibit sharp transition if the accuracy is defined as the fraction of points whose regression loss is smaller than a small threshold. All these works are limited to linear models, while our work is able to analyze deep homogeneous neural nets.

Concurrent Works. Concurrent to our work, [Kumar et al. \(2024\)](#) hypothesized that grokking is due to a transition from lazy to rich regimes, and observed that manipulating the scale parameter for the output and the task-model alignment between NTK and ground-truth labels can control the time to escape the lazy regime and the test accuracy when staying in the lazy regime. This observation is consistent with our work, but we provide rigorous analyses for both lazy and rich regimes with quantitative bounds. [Xu et al. \(2024\)](#) proved that training two-layer neural nets on XOR cluster data with noisy labels can exhibit grokking. Both their work and our work go beyond linear models, but their analysis does not show whether the transition in test accuracy is sharp or not, while our work provides quantitative bounds for the sharp transition in grokking.

Implicit Bias. A line of works seek to characterize the implicit bias of optimization algorithms that drive them to generalizable solutions. Several forms of implicit bias have been considered, including equivalence to NTK (Du et al., 2019b;a; Allen-Zhu et al., 2019; Zou et al., 2020; Chizat et al., 2019; Arora et al., 2019b; Ji and Telgarsky, 2020b; Cao and Gu, 2019), margin maximization (Soudry et al., 2018; Nacson et al., 2019a; Lyu and Li, 2020; Ji and Telgarsky, 2020a), parameter norm minimization (Gunasekar et al., 2017; 2018a; Arora et al., 2019a) and sharpness minimization (Blanc et al., 2020; Damian et al., 2021; HaoChen et al., 2021; Li et al., 2022; Lyu et al., 2022; Gu et al., 2023). In this work, we characterize the early phase implicit bias based on NTK and late phase implicit bias based on margin maximization.

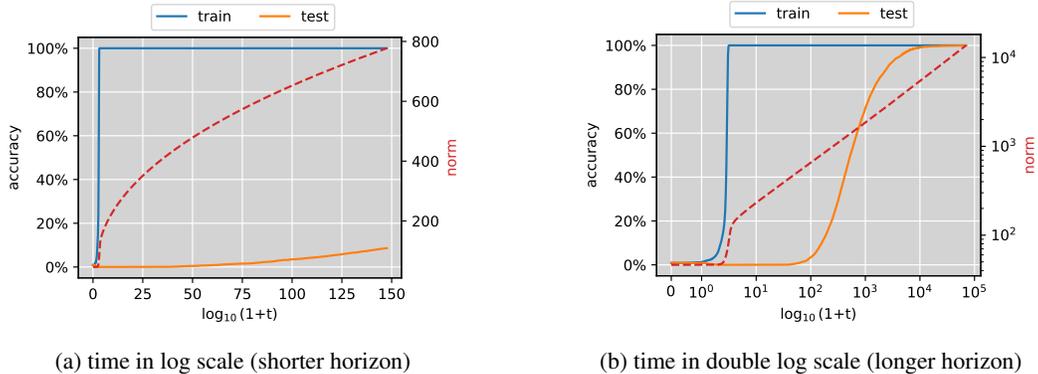


Figure 4: Training two-layer ReLU nets for modular addition exhibits grokking without weight decay, but the transition is no longer sharp: it can take exponentially longer than the time needed for reaching 100% training accuracy.

B DISCUSSION: OTHER IMPLICIT BIASES CAN ALSO INDUCE GROKKING

The main paper focuses on the grokking phenomenon induced by large initialization and small weight decay, where the transition in the implicit bias is very sharp: $\frac{1-c}{\lambda} \log \alpha$ leads to the kernel predictor, but increasing it slightly to $\frac{1+c}{\lambda} \log \alpha$ leads to a KKT solution of min-norm/max-margin problems associated with the neural net. In this section, we discuss two other possible sources of grokking: implicit margin maximization and sharpness reduction. These two sources can also cause the neural net to generalize eventually, but the transition in test accuracy is not as sharp as the case of using large initialization and small weight decay.

B.1 IMPLICIT MARGIN MAXIMIZATION WITHOUT WEIGHT DECAY

One may wonder if grokking still exists if we remove the small weight decay in our setting. In Liu et al. (2023), it has been argued that adding a non-zero weight decay is crucial for grokking since the weight norm has to lie in a certain range (Goldilocks zone) to generalize.

We note that for classification tasks, Lyu and Li (2020); Ji and Telgarsky (2020a) have shown that gradient descent/flow on homogeneous neural nets converges to KKT points of the margin maximization problem (R1), even without weight decay. So it is still possible that some implicit bias traps the solution in early phase, such as the implicit bias to NTK Allen-Zhu et al. (2019); ?; Ji and Telgarsky (2020b); Cao and Gu (2019). Then after training for sufficiently long time, the margin maximization bias arises and leads to grokking.

We empirically show that training two-layer neural nets without weight decay indeed leads to grokking, but the transition is no longer sharp. More specifically, we use the same setting as in Section 2 but with weight decay $\lambda = 0$. We focus on analyzing the training dynamics of full-batch gradient descent with learning rate so small that the trajectory is close to the gradient flow, which ensures that gradient noise/training instability is not a potential source of implicit regularization.

However, directly training with a constant learning rate leads to very slow loss convergence in the late phase. Inspired by Lyu and Li (2020); Nacson et al. (2019b), we set the learning rate as $\frac{\eta}{\mathcal{L}(\theta(t))}$, where η is a constant. The main intuition behind this is that the Hessian of $\mathcal{L}(\theta)$ can be shown to be bounded by $\mathcal{O}(\mathcal{L}(\theta) \cdot \text{poly}(\|\theta\|_2))$, and gradient descent stays close to gradient flow when the learning rate is much lower than the reciprocal of the smoothness. Figure 4 plots training and test curves, where x -axis shows the corresponding continuous time calculated by summing the learning rates up to this point. It is shown that a very quick convergence in training accuracy occurs in the early phase, but the test accuracy improves extremely slowly in the late phase: the test accuracy slowly transitions from $\sim 0\%$ to 100% as t increases from 10^{10^2} to 10^{10^4} . In contrast, Figure 2 shows that training with weight decay leads to a sharp transition when the number of steps is $\sim 10^7$.

We note that this much longer time for grokking is not surprising, since the convergence to max-margin solutions is already as slow as $\mathcal{O}(\frac{\log \log t}{\log t})$ for linear logistic regression (Soudry et al., 2018).

B.2 SHARPNESS REDUCTION WITH LABEL NOISE

Grokking can also be triggered by the sharpness-reduction implicit bias of optimization algorithm, e.g., label noise SGD (Blanc et al., 2020; Damian et al., 2021; Li et al., 2022) and 1-SAM (Wen et al., 2023). The high-level idea here is that the minimizers of the training loss connect as a Riemannian manifold due to the overparametrization in the model. Training usually contains two phases for algorithms which implicitly minimize the sharpness when the learning rate is small, where in the first phase the algorithm minimizes the training loss and finds a minimizer with poor generalization, while in the second phase the algorithm implicitly minimizes the sharpness along the manifold of the minimizers, and the model groks when the sharpness is sufficiently small. The second phase is typically longer than the first phase by magnitude and this causes grokking.

As a concrete example, Li et al. (2022) shows that grokking happens when training two-layer diagonal linear networks (4) by label noise SGD (6) when the learning rate is sufficiently small. We consider the same initialization as in Section 3.2.3, where $\theta(0) = \alpha(\mathbf{1}, \mathbf{1})$, where $\alpha > 0$ controls the initialization scale. Different from Section 3.2.3, we consider a ℓ_2 loss, and the label is generated by a sparse linear ground-truth, that is, $\mathbf{y}^*(\mathbf{x}) = \mathbf{x}^\top \mathbf{w}^*$, where \mathbf{w}^* is κ -sparse. We also assume data distribution is the isotropic Gaussian distribution $\mathcal{N}(0, I)$.

$$\theta(t+1) = \theta(t) - \eta \nabla_{\theta} (f(\theta; \mathbf{x}_{i_t}) - y_{i_t} + \xi_t)^2, \text{ where } \xi_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1) \quad (6)$$

Theorem B.1 *There exists $C_1 > 0, C_2 \in (0, 1)$, such that for any n, d satisfying $C_1 \kappa \ln d \leq n \leq d(1 - C_2)$, the following holds with $1 - e^{-\Omega(n)}$ probability with respect to randomness of training data:*

1. *The gradient flow starting of \mathcal{L} from $\theta(0)$ has a limit, θ_{GF} , which achieves zero training loss. The expected population loss θ_{GF} is at least $\Omega(\|\mathbf{w}^*\|_2^2)$. (Theorem 6.7)*
2. *For any $\alpha \in (1, 2)$ and $T > 0$, $\theta(T/\eta^\alpha)$ converges to θ_{GF} in distribution as $\eta \rightarrow 0$. (Implication of Theorem B.9)*
3. *For any $T > 0$, $\theta(T/\eta^2)$ converges to some $\bar{\theta}(T)$ in distribution as $\eta \rightarrow 0$ and $\bar{\theta}(T)$ has zero training loss. Moreover, the population loss of $\bar{\theta}(T)$ converges to 0 as $T \rightarrow \infty$. (Theorem 6.1)*

The above theorem shows that a sharp transition of the implicit biases can be seen in the log time scale: when $\log t = \log T + \alpha \log \frac{1}{\eta}$ for any $\alpha \in (1, 2)$, it converges to θ_{GF} ; but increasing the time slightly in the log scale to $\log t = \log T + 2 \log \frac{1}{\eta}$ leads to $\bar{\theta}(T)$, which minimizes the population loss to 0. However, this transition is less sharp than our main results, where we show that increasing the time slightly in the normal scale, from $\frac{1-c}{\lambda} \log \alpha$ to $\frac{1+c}{\lambda} \log \alpha$, changes the implicit bias.

C PROOFS FOR CLASSIFICATION WITH L2 REGULARIZATION

C.1 PRELIMINARIES

Let $F_{\text{LL}} : [1, +\infty) \rightarrow [0, +\infty)$, $x \mapsto x \log x$ be the linear-times-log function, and $F_{\text{LL}}^{-1} : [0, +\infty) \rightarrow [1, +\infty)$ be its inverse function.

Lemma C.1 For all $y > 0$, $F_{\text{LL}}^{-1}(y) \geq \frac{y}{\log(y+1)}$.

Proof: Since F_{LL} is an increasing function, it suffices to show $F_{\text{LL}}(\frac{y}{\log(y+1)}) \leq y$, and this can be indeed proved by

$$F_{\text{LL}}\left(\frac{y}{\log(y+1)}\right) \leq \frac{y}{\log(y+1)} \log \frac{y}{\log(y+1)} = y \cdot \frac{\log \frac{y}{\log(y+1)}}{\log(y+1)} \leq y,$$

where the last inequality holds because $\frac{y}{\log(y+1)} \leq y+1$. \square

Lemma C.2 For all $y > 0$ and $\beta \in (0, 1)$,

$$F_{\text{LL}}^{-1}(\beta y) - 1 \geq \beta(F_{\text{LL}}^{-1}(y) - 1).$$

Proof: By concavity of F_{LL}^{-1} , $F_{\text{LL}}^{-1}(\beta y) \geq \beta F_{\text{LL}}^{-1}(y) + (1 - \beta)F_{\text{LL}}^{-1}(0)$. Rearranging the terms gives the desired result. \square

Lemma C.3 The following holds for $r_1, \dots, r_n > 0$ and $\mathcal{L} = \frac{1}{n} \sum_{i=1}^n r_i$:

$$\mathcal{L} \log \frac{1}{n\mathcal{L}} \leq \frac{1}{n} \sum_{i=1}^n r_i \log \frac{1}{r_i} \leq \mathcal{L} \log \frac{1}{\mathcal{L}}.$$

Proof: For proving the first inequality, note that $\log \frac{1}{r_i} \geq \log \frac{1}{n\mathcal{L}}$, so

$$\frac{1}{n} \sum_{i=1}^n r_i \log \frac{1}{r_i} \geq \frac{1}{n} \sum_{i=1}^n r_i \log \frac{1}{n\mathcal{L}} = \mathcal{L} \log \frac{1}{n\mathcal{L}}.$$

For the other inequality, by Jensen's inequality, we have

$$\frac{1}{n} \sum_{i=1}^n r_i \log \frac{1}{r_i} = -\frac{1}{n} \sum_{i=1}^n F_{\text{LL}}(r_i) \leq -F_{\text{LL}}(\mathcal{L}) = \mathcal{L} \log \frac{1}{\mathcal{L}},$$

which completes the proof. \square

C.2 KERNEL REGIME

Now we present the proof of [Theorem 3.4](#) for the kernel regime. We follow the notations in [Section 3.2.1](#) and define $\mathbf{h}_{\text{ntk}}^*$ as the unique unit vector along the direction of the optimal solution to [\(K1\)](#) and $\gamma_{\text{ntk}} := \min_{i \in [n]} \{y_i \langle \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle\}$. The kernel regime is a training regime where the parameter does not move very far from the initialization. Mathematically, let $\epsilon_{\text{max}} > 0$ be a small constant so that

$$\max_{i \in [n]} \{\|\nabla f_i(\boldsymbol{\theta}) - \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 : i \in [n]\} < \frac{\gamma_{\text{ntk}}}{2}$$

holds for all $\|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 < \epsilon_{\text{max}}$, and let $T_{\text{max}} := \inf \left\{ t \geq 0 : \left\| \frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 > \epsilon_{\text{max}} \right\}$ be the time that $\boldsymbol{\theta}(t)$ moves far enough to change the NTK features significantly.

A key difference to existing analyses of the kernel regime is that our setting has weight decay, so the norm can change a lot even when the parameter direction has not changed much yet. Therefore, we normalize the parameter by multiplying a factor of $\frac{e^{\lambda t}}{\alpha}$ and compare it to the initial parameter direction. We consider $\boldsymbol{\theta}(t)$ as moving too far only if $\left\| \frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2$ is too large, which is different from existing analyses in that they usually consider the absolute moving distance $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\|_2$.

In the following, we first derive an upper bound for the loss convergence rate, and then use it to lower bound the training time in the kernel regime $T_{\max} \geq \frac{1}{\lambda}(\log \alpha - \frac{1}{L} \log \log A + \Omega(1))$. Finally, we establish the implicit bias towards the kernel predictor with a more careful analysis of the trajectory.

More notations are needed for the proof. Let $q_i(\boldsymbol{\theta}) := y_i f_i(\boldsymbol{\theta})$ be the output margin of the neural net at the i -th training sample. Let $r_i(\boldsymbol{\theta}) = \ell(f_i(\boldsymbol{\theta}), y_i) = \exp(-q_i(\boldsymbol{\theta}))$ be the loss incurred at the i -th training sample. Let $A := \frac{\alpha^{2(L-1)}}{\lambda}$, which is a factor that will appear very often in our proof. Let $G_{\text{ntk}} := \sup\{\|\nabla f_i(\boldsymbol{\theta})\|_2 : \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 < \epsilon_{\max}, i \in [n]\}$ be the maximum norm of NTK features. The following lemma gives an upper bound for the training loss over time.

Lemma C.4 For all $0 \leq t \leq T_{\max}$,

$$\frac{d\mathcal{L}}{dt} \leq -\frac{\gamma_{\text{ntk}}^2}{4} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \mathcal{L}^2 + \lambda L \cdot \mathcal{L} \log \frac{1}{\mathcal{L}}, \quad (7)$$

$$\mathcal{L}(\boldsymbol{\theta}(t)) \leq \max \left\{ \frac{1}{1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A (1 - e^{-2(L-1)\lambda t})}, \frac{e^{2(L-1)\lambda t}}{F_{\text{LL}}^{-1}(\frac{\gamma_{\text{ntk}}^2}{8L} A) - 1} \right\}. \quad (8)$$

Proof: By the update rule of $\boldsymbol{\theta}$, the following holds for all $t > 0$:

$$\frac{d\mathcal{L}}{dt} = \left\langle \nabla \mathcal{L}(\boldsymbol{\theta}), \frac{d\boldsymbol{\theta}}{dt} \right\rangle = - \underbrace{\left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla q_i(\boldsymbol{\theta}) \right\|_2^2}_{=: V_1} + \underbrace{\frac{\lambda L}{n} \sum_{i=1}^n r_i q_i}_{=: V_2}.$$

For V_1 , we have

$$-V_1 = \left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla q_i(\boldsymbol{\theta}) \right\|_2^2 \geq \left(\frac{1}{n} \sum_{i=1}^n r_i \langle \nabla q_i(\boldsymbol{\theta}), \mathbf{h}_{\text{ntk}}^* \rangle \right)^2. \quad (9)$$

Note that for all $t \leq T_{\max}$,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n r_i \langle \nabla q_i(\boldsymbol{\theta}), \mathbf{h}_{\text{ntk}}^* \rangle &\geq \frac{1}{n} \sum_{i=1}^n r_i (\langle \nabla q_i(\alpha e^{-\lambda t} \bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle - \|\nabla f_i(\boldsymbol{\theta}) - \nabla f_i(\alpha e^{-\lambda t} \bar{\boldsymbol{\theta}}_{\text{init}})\|_2) \\ &\geq \frac{1}{n} \sum_{i=1}^n r_i \cdot \alpha^{L-1} e^{-(L-1)\lambda t} \left(\gamma_{\text{ntk}} - \frac{1}{2} \gamma_{\text{ntk}} \right) \\ &= \frac{\gamma_{\text{ntk}}}{2} \alpha^{L-1} e^{-(L-1)\lambda t} \mathcal{L}, \end{aligned}$$

where the second inequality is due to $\langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle \geq \gamma_{\text{ntk}}$, $\left\| \nabla f_i(\frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta}(t)) - \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}) \right\|_2 \leq \frac{\gamma_{\text{ntk}}}{2}$ and the $(L-1)$ -homogeneity of ∇f_i . Together with (9), we obtain the following upper bound for V_1 :

$$V_1 \leq -\frac{\gamma_{\text{ntk}}^2}{4} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \mathcal{L}^2. \quad (10)$$

For V_2 , by Lemma C.3 we have $\frac{1}{n} \sum_{i=1}^n r_i \log \frac{1}{r_i} \leq \mathcal{L} \log \frac{1}{\mathcal{L}}$. So we have the following upper bound for V_2 :

$$V_2 \leq \lambda L \cdot \mathcal{L} \log \frac{1}{\mathcal{L}}. \quad (11)$$

Combining this with (10) together proves (7).

For proving (8), we first divide by \mathcal{L}^2 on both sides of (7) to get

$$\frac{d}{dt} \frac{1}{\mathcal{L}} \geq \frac{\gamma_{\text{ntk}}^2}{4} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} - \lambda L \cdot \frac{1}{\mathcal{L}} \log \frac{1}{\mathcal{L}}. \quad (12)$$

Let $\mathcal{E} := \{t \in [0, T_{\max}] : F_{LL}(\frac{1}{\mathcal{L}}) \geq \frac{\gamma_{\text{ntk}}^2}{8L} A e^{-2(L-1)\lambda t}\}$. For all $t \in [0, T_{\max}]$, if $t \in \mathcal{E}$, then by [Lemma C.2](#),

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}(t)) &\leq \frac{1}{F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A e^{-2(L-1)\lambda t}\right)} \leq \frac{1}{1 + \left(F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A\right) - 1\right) e^{-2(L-1)\lambda t}} \\ &\leq \frac{1}{\left(F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A\right) - 1\right) e^{-2(L-1)\lambda t}} \leq \frac{e^{2(L-1)\lambda t}}{F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A\right) - 1}. \end{aligned}$$

Otherwise, let t' be the largest number in $\mathcal{E} \cup \{0\}$ that is smaller than t . Then

$$\begin{aligned} \frac{1}{\mathcal{L}(\boldsymbol{\theta}(t))} &= \frac{1}{\mathcal{L}(\boldsymbol{\theta}(t'))} + \int_{t'}^t \frac{d}{d\tau} \frac{1}{\mathcal{L}} d\tau \\ &\geq \frac{1}{\mathcal{L}(\boldsymbol{\theta}(t'))} + \int_{t'}^t \frac{\gamma_{\text{ntk}}^2}{8} \alpha^{2(L-1)} e^{-2(L-1)\lambda\tau} d\tau \\ &= \frac{1}{\mathcal{L}(\boldsymbol{\theta}(t'))} + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A (e^{-2(L-1)\lambda t'} - e^{-2(L-1)\lambda t}). \end{aligned}$$

If $t' \in \mathcal{E}$, then $\frac{1}{\mathcal{L}(\boldsymbol{\theta}(t'))} \geq F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A e^{-2(L-1)\lambda t'}\right) \geq F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A e^{-2(L-1)\lambda t}\right)$, which contradicts to $t \notin \mathcal{E}$. So it must hold that $t' = 0$, and thus $\frac{1}{\mathcal{L}(\boldsymbol{\theta}(t))} \geq 1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A (1 - e^{-2(L-1)\lambda t})$. Combining all these together proves [\(8\)](#). \square

Lemma C.5 *There exists a constant ΔT such that $T_{\max} \geq \frac{1}{\lambda} (\log \alpha - \frac{1}{L} \log \log A + \Delta T)$ when α is sufficiently large. Furthermore, for all $t \leq \frac{1}{\lambda} (\log \alpha - \frac{1}{L} \log \log A + \Delta T)$,*

$$\left\| \frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 = \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}.$$

Proof: By definition of T_{\max} , to prove the first claim, it suffices to show that there exists a constant ΔT such that $\left\| \frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 \leq \epsilon_{\max}$ for all $t \leq \min\{\frac{1}{\lambda} (\log \alpha - \frac{1}{L} \log \log A + \Delta T), T_{\max}\}$.

When $t \leq T_{\max}$, we have the following gradient upper bound:

$$\begin{aligned} \|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_2 &= \left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla q_i(\boldsymbol{\theta}(t)) \right\|_2 \leq \frac{1}{n} \sum_{i=1}^n r_i \|\nabla q_i(\boldsymbol{\theta}(t))\|_2 \\ &\leq \frac{1}{n} \sum_{i=1}^n r_i \alpha^{L-1} e^{-(L-1)\lambda t} \|\nabla q_i(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t))\|_2 \\ &\leq G_{\text{ntk}} \alpha^{L-1} e^{-(L-1)\lambda t} \mathcal{L}. \end{aligned}$$

By update rule, $e^{\lambda t} \boldsymbol{\theta}(t) - \alpha \bar{\boldsymbol{\theta}}_{\text{init}} = \int_0^t e^{\lambda\tau} \nabla \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau$. Applying the gradient upper bounds gives

$$\begin{aligned} \frac{1}{\alpha} \left\| e^{\lambda t} \boldsymbol{\theta}(t) - \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 &\leq \frac{1}{\alpha} \int_0^t \left\| e^{\lambda\tau} \nabla \mathcal{L}(\boldsymbol{\theta}(\tau)) \right\|_2 d\tau \\ &\leq \frac{1}{\alpha} \int_0^t e^{\lambda\tau} G_{\text{ntk}} \alpha^{L-1} e^{-(L-1)\lambda\tau} \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau \\ &= G_{\text{ntk}} \alpha^{L-2} \int_0^t e^{-(L-2)\lambda\tau} \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau. \end{aligned}$$

Substituting \mathcal{L} in $\int_0^t e^{-(L-2)\lambda\tau} \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau$ with the loss upper bound in [Lemma C.4](#), we obtain the following for all $t \leq T_{\max}$,

$$\int_0^t e^{-(L-2)\lambda\tau} \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau \leq \int_0^t \max \left\{ \frac{e^{-(L-2)\lambda\tau}}{1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A (1 - e^{-2(L-1)\lambda\tau})}, \frac{e^{-(L-2)\lambda\tau} \cdot e^{2(L-1)\lambda\tau}}{F_{LL}^{-1}\left(\frac{\gamma_{\text{ntk}}^2}{8L} A\right) - 1} \right\} d\tau.$$

So we have $\int_0^t e^{-(L-2)\lambda\tau} \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau \leq \max\{I_1, I_2\}$, where

$$I_1 := \int_0^t \frac{d\tau}{1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A (1 - e^{-2(L-1)\lambda\tau})}, \quad I_2 := \int_0^t \frac{e^{L\lambda\tau}}{F_{\text{LL}}^{-1}(\frac{\gamma_{\text{ntk}}^2}{8L} A) - 1} d\tau.$$

We can exactly compute the integrals for both I_1 and I_2 . For I_1 , it holds for all $t \leq \min\{T_{\text{max}}, \mathcal{O}(\frac{1}{\lambda} \log \alpha)\}$ that

$$\begin{aligned} I_1 &= \frac{1}{2(L-1)\lambda \left(1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A\right)} \log \left(\left(1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A\right) e^{2(L-1)\lambda t} - \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A \right) \\ &\leq \frac{1}{2(L-1)\lambda \left(1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A\right)} \log \left(\left(1 + \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A\right) \text{poly}(\alpha) - \frac{\gamma_{\text{ntk}}^2}{16(L-1)} A \right) \\ &= \mathcal{O}\left(\frac{\log A}{\lambda A}\right). \end{aligned}$$

For I_2 , it holds for all $t \leq T_{\text{max}}$ that

$$I_2 = \frac{e^{L\lambda t} - 1}{L\lambda \cdot \left(F_{\text{LL}}^{-1}(\frac{\gamma_{\text{ntk}}^2}{8L} A) - 1\right)} = \mathcal{O}\left(\frac{\log A}{\lambda A}\right) \cdot e^{L\lambda t}.$$

Putting all these together, we have

$$\begin{aligned} \frac{1}{\alpha} \|e^{\lambda t} \boldsymbol{\theta}(t) - \alpha \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 &= \mathcal{O}(\alpha^{L-2}) \cdot \left(\mathcal{O}\left(\frac{\log A}{\lambda A}\right) + \mathcal{O}\left(\frac{\log A}{\lambda A}\right) \cdot e^{L\lambda t} \right) \\ &= \mathcal{O}\left(\frac{\alpha^{L-2} \log A}{\lambda A}\right) \cdot e^{L\lambda t} = \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}. \end{aligned} \quad (13)$$

Therefore, there exists a constant $C > 0$ such that when α is sufficiently large, $\|\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 \leq C\alpha^{-L} \log A \cdot e^{L\lambda t}$ for all $t \leq \min\{T_{\text{max}}, \mathcal{O}(\frac{1}{\lambda} \log \alpha)\}$. Setting $\Delta T := \frac{1}{L} \log \frac{\epsilon_{\text{max}}}{C}$, we obtain the following bound for all $t \leq \min\{T_{\text{max}}, \frac{1}{\lambda} (\log \alpha - \frac{1}{L} \log \log A + \Delta T)\}$:

$$\left\| \frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 \leq C\alpha^{-L} \log A \cdot \exp\left(L \log \alpha - \log \log A + \log \frac{\epsilon_{\text{max}}}{C}\right) = \epsilon_{\text{max}},$$

which implies $T_{\text{max}} \geq \frac{1}{\lambda} (\log \alpha - \frac{1}{L} \log \log A + \Delta T)$. Combining this with (13) proves the second claim. \square

Now we turn to analyze the implicit bias. Let $\boldsymbol{\delta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) := (\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) - \bar{\boldsymbol{\theta}}_{\text{init}})$ and $\mathbf{h}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) := \alpha^L e^{-L\lambda t} \boldsymbol{\delta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})$. The following lemma shows that the output function can be approximated by a linear function.

Lemma C.6 For all $t \leq \frac{1}{\lambda} (\log \alpha - \frac{1}{L} \log \log A + \Delta T)$,

$$q_i(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}) = \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \boldsymbol{\delta} \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \|\boldsymbol{\delta}\|_2, \quad (14)$$

$$q_i(\boldsymbol{\theta}) = \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h} \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \|\mathbf{h}\|_2. \quad (15)$$

Proof: By Lemma C.5 and Taylor expansion, we have

$$\begin{aligned} q_i(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}) &= q_i(\bar{\boldsymbol{\theta}}_{\text{init}}) + \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \boldsymbol{\delta} \rangle + \mathcal{O}(\|\boldsymbol{\delta}\|_2^2) \\ &= q_i(\bar{\boldsymbol{\theta}}_{\text{init}}) + \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \boldsymbol{\delta} \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \|\boldsymbol{\delta}\|_2 \\ &= \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \boldsymbol{\delta} \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \|\boldsymbol{\delta}\|_2, \end{aligned}$$

which proves (14). Combining this with the L -homogeneity of f_i proves (15). \square

In the following two lemmas, we derive lower bounds for the loss convergence that will be used in the implicit bias analysis later.

Lemma C.7 For all $0 \leq t \leq T_{\max}$,

$$\frac{d\mathcal{L}}{dt} \geq -G_{\text{ntk}}^2 \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \mathcal{L}^2 + \lambda L \cdot \mathcal{L} \log \frac{1}{n\mathcal{L}}, \quad (16)$$

$$\mathcal{L}(\theta(t)) \geq \min \left\{ \frac{n}{F_{\text{LL}}^{-1} \left(\frac{2}{nL} G_{\text{ntk}}^2 A e^{-2(L-1)\lambda t} \right)}, \frac{1}{n} \exp(1 - 2(L-1)\lambda) \right\}. \quad (17)$$

Proof: The proof is similar to that of [Lemma C.4](#), but now we are proving lower bounds. By the update rule of θ , the following holds for all $t > 0$:

$$\frac{d\mathcal{L}}{dt} = \left\langle \nabla \mathcal{L}(\theta), \frac{d\theta}{dt} \right\rangle = - \underbrace{\left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla q_i(\theta) \right\|_2^2}_{=:V_1} + \underbrace{\frac{\lambda L}{n} \sum_{i=1}^n r_i \log \frac{1}{r_i}}_{=:V_2}.$$

For V_1 , we can lower bound it as follows:

$$\begin{aligned} V_1 &\geq - \left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla q_i(\theta) \right\|_2^2 \\ &\geq -\alpha^{2(L-1)} e^{-2(L-1)\lambda t} \left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \theta\right) \right\|_2^2 \\ &\geq -\alpha^{2(L-1)} e^{-2(L-1)\lambda t} G_{\text{ntk}}^2 \mathcal{L}^2, \end{aligned}$$

where the last inequality holds because $\|\nabla q_i(\frac{1}{\alpha} e^{\lambda t} \theta)\|_2 \leq \|\nabla q_i(\bar{\theta}_{\text{init}})\|_2 + \|\nabla q_i(\frac{1}{\alpha} e^{\lambda t} \theta) - \nabla q_i(\bar{\theta}_{\text{init}})\|_2 \leq G_{\text{ntk}}$.

For V_2 , by [Lemma C.3](#) we have $\frac{1}{n} \sum_{i=1}^n r_i \log \frac{1}{r_i} \geq \mathcal{L} \log \frac{1}{n\mathcal{L}}$. So $V_2 \geq \lambda L \cdot \mathcal{L} \log \frac{1}{n\mathcal{L}}$. Combining the inequalities for V_1 and V_2 together proves (16).

For proving (8), we first divide by \mathcal{L}^2 on both sides of (16) to get

$$\frac{d}{dt} \frac{1}{\mathcal{L}} \leq G_{\text{ntk}}^2 \alpha^{2(L-1)} e^{-2(L-1)\lambda t} - \lambda L \cdot \frac{1}{\mathcal{L}} \log \frac{1}{n\mathcal{L}}. \quad (18)$$

Let $\mathcal{E} := \{t \in [0, T_{\max}] : F_{\text{LL}}(\frac{1}{n\mathcal{L}}) \leq \frac{2}{nL} G_{\text{ntk}}^2 A e^{-2(L-1)\lambda t}\}$. It is easy to see that $0 \in \mathcal{E}$. For all $t \in [0, T_{\max}]$, if $t \in \mathcal{E}$, then by [Lemma C.2](#),

$$\mathcal{L}(\theta(t)) \geq \frac{n}{F_{\text{LL}}^{-1} \left(\frac{2}{nL} G_{\text{ntk}}^2 A e^{-2(L-1)\lambda t} \right)}.$$

If no t is at the boundary of \mathcal{E} , then we are done. Otherwise, let $t > 0$ be one on the boundary of \mathcal{E} . Then at time t ,

$$\frac{d}{dt} \frac{1}{n\mathcal{L}} = -\frac{\lambda L}{2} \cdot \frac{1}{n\mathcal{L}} \log \frac{1}{n\mathcal{L}}.$$

So as long as $1 + \log \frac{1}{n\mathcal{L}(\theta(t))} > 2(L-1)\lambda$,

$$\begin{aligned} \frac{d}{dt} F_{\text{LL}}\left(\frac{1}{n\mathcal{L}}\right) &= \left(1 + \log \frac{1}{n\mathcal{L}}\right) \frac{d}{dt} \frac{1}{n\mathcal{L}} \\ &= -\left(1 + \log \frac{1}{n\mathcal{L}}\right) \cdot \frac{\lambda L}{2} \cdot \frac{1}{n\mathcal{L}} \log \frac{1}{n\mathcal{L}} \\ &= -\left(1 + \log \frac{1}{n\mathcal{L}}\right) \cdot \frac{2}{nL} G_{\text{ntk}}^2 A e^{-2(L-1)\lambda t} < \left(\frac{2}{nL} G_{\text{ntk}}^2 A e^{-2(L-1)\lambda t}\right)', \end{aligned}$$

which means t is not actually on the boundary. So $1 + \log \frac{1}{n\mathcal{L}(\theta(t))} \leq 2(L-1)\lambda$. Repeating this argument with inequalities and taking some integrals can show that every point outside \mathcal{E} satisfies $1 + \log \frac{1}{n\mathcal{L}(\theta(t))} \leq 2(L-1)\lambda$. \square

Lemma C.8 For any constant $c \in (0, 1)$, letting $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$, it holds that

$$\frac{d\|\mathbf{h}\|_2}{dt} = \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \left\langle \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right), \frac{\mathbf{h}}{\|\mathbf{h}\|_2} \right\rangle - L\lambda \|\mathbf{h}\|_2, \quad (19)$$

$$\frac{d}{dt} \log \frac{1}{\mathcal{L}} \geq (\gamma_{\text{ntk}} + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}) \frac{d\|\mathbf{h}\|_2}{dt}. \quad (20)$$

Proof: By update rule,

$$\begin{aligned} \frac{d\mathbf{h}}{dt} &= \alpha^L e^{-L\lambda t} \frac{d\boldsymbol{\delta}}{dt} - L\lambda \alpha^L e^{-L\lambda t} \boldsymbol{\delta} \\ &= \alpha^L e^{-L\lambda t} \left(\frac{e^{\lambda t}}{\alpha} \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}(t)) \nabla q_i(\boldsymbol{\theta}(t)) \right) - L\lambda \mathbf{h} \\ &= \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}(t)) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t)\right) - L\lambda \mathbf{h}. \end{aligned}$$

Then we can deduce (19) from $\frac{d}{dt} \|\mathbf{h}\|_2 = \left\langle \frac{d\mathbf{h}}{dt}, \frac{\mathbf{h}}{\|\mathbf{h}\|_2} \right\rangle$.

For proving (20), first, we apply the chain rule to get

$$\frac{d}{dt} \log \frac{1}{\mathcal{L}} = \frac{1}{\mathcal{L}(\boldsymbol{\theta})} \langle -\nabla \mathcal{L}(\boldsymbol{\theta}) - \lambda \boldsymbol{\theta}, -\nabla \mathcal{L}(\boldsymbol{\theta}) \rangle = \underbrace{\frac{1}{\mathcal{L}(\boldsymbol{\theta})} \|\nabla \mathcal{L}(\boldsymbol{\theta})\|_2^2}_{=:V_1} - \underbrace{\frac{\lambda}{\mathcal{L}(\boldsymbol{\theta})} \langle \boldsymbol{\theta}, -\nabla \mathcal{L}(\boldsymbol{\theta}) \rangle}_{=:V_2}.$$

For V_1 , note that

$$\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_2 = \left\| \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i(\boldsymbol{\theta}) \right\|_2 = \alpha^{L-1} e^{-(L-1)\lambda t} \underbrace{\left\| \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right) \right\|_2}_{=:V_3}.$$

By Lemma C.5, $\|\boldsymbol{\delta}(t)\|_2 = \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}$, so $\langle \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right), \mathbf{h}_{\text{ntk}}^* \rangle = \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}$. One way to give a lower bound for V_3 is to consider the projection of $\frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right)$ along $\mathbf{h}_{\text{ntk}}^*$:

$$\begin{aligned} V_3 &\geq \left\langle \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right), \mathbf{h}_{\text{ntk}}^* \right\rangle \\ &= \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) (\langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}) \\ &\geq \mathcal{L}(\boldsymbol{\theta}) \cdot (\gamma_{\text{ntk}} + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}) > 0. \end{aligned}$$

Another way to give a lower bound for V_3 is to consider the projection of $\frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right)$ along \mathbf{h} :

$$V_3 \geq \left\langle \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right), \frac{\mathbf{h}}{\|\mathbf{h}\|_2} \right\rangle.$$

Putting these two lower bounds together gives:

$$\begin{aligned} V_1 &= \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \frac{V_3^2}{\mathcal{L}(\boldsymbol{\theta})} \\ &\geq (\gamma_{\text{ntk}} + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}) \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \left\langle \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \nabla q_i\left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}\right), \frac{\mathbf{h}}{\|\mathbf{h}\|_2} \right\rangle. \end{aligned}$$

For V_2 , by [Lemma C.6](#), we have

$$\begin{aligned}
V_2 &= \frac{\lambda}{\mathcal{L}(\boldsymbol{\theta})} \cdot \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}) \langle \boldsymbol{\theta}, \nabla q_i(\boldsymbol{\theta}) \rangle \\
&= L\lambda \cdot \frac{1}{n\mathcal{L}(\boldsymbol{\theta})} \sum_{i=1}^n r_i(\boldsymbol{\theta}) q_i(\boldsymbol{\theta}) \\
&= L\lambda \cdot \frac{1}{n\mathcal{L}(\boldsymbol{\theta})} \sum_{i=1}^n r_i(\boldsymbol{\theta}) (\langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h} \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \|\mathbf{h}\|_2) \\
&\leq L\lambda \|\mathbf{h}\|_2 (\gamma_{\text{ntk}} + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t}).
\end{aligned}$$

Putting the bounds for V_1 and V_2 together proves [\(20\)](#). \square

Now we are ready to prove the directional convergence of $\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})$ to $\mathbf{h}_{\text{ntk}}^*$.

Theorem C.9 For any constant $c \in (0, 1)$, letting $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$, it holds that

$$\lim_{\alpha \rightarrow +\infty} \frac{\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})}{\frac{1}{\gamma_{\text{ntk}}} \log \frac{\alpha^c}{\lambda}} = \mathbf{h}_{\text{ntk}}^*.$$

Proof: Integrating [\(20\)](#) in [Lemma C.8](#) over $t \in [0, T_c^-(\alpha)]$ gives

$$\log \frac{1}{\mathcal{L}(\boldsymbol{\theta}(T_c^-(\alpha)))} \geq (\gamma_{\text{ntk}} + \mathcal{O}(\alpha^{-cL} \log A)) \|\mathbf{h}(T_c^-(\alpha))\|_2.$$

By [\(8\)](#) in [Lemma C.4](#) and [\(17\)](#) in [Lemma C.7](#), we have

$$\begin{aligned}
\log \frac{1}{\mathcal{L}(\boldsymbol{\theta}(T_c^-(\alpha)))} &= \log(F_{\text{LL}}^{-1}(\Theta(A)) - 1) - 2(L-1)\lambda T_c^-(\alpha) \\
&= \log A - 2(L-1)(1-c) \log \alpha + O(\log \log A) \\
&= \log \frac{\alpha^c}{\lambda} + O(\log \log A) \rightarrow +\infty.
\end{aligned}$$

By [Lemma C.6](#), $\log \frac{1}{\mathcal{L}(\boldsymbol{\theta}(T_c^-(\alpha)))} \leq O(\|\mathbf{h}(T_c^-(\alpha))\|_2)$, so we also have $\|\mathbf{h}(T_c^-(\alpha))\|_2 \rightarrow +\infty$.

Then

$$\begin{aligned}
\gamma_{\text{ntk}} &\geq \frac{\min_{i \in [n]} \langle \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h} \rangle}{\|\mathbf{h}\|_2} \geq \frac{\min_{i \in [n]} \{q_i(\boldsymbol{\theta})\}}{\|\mathbf{h}\|_2} + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda T_c^-(\alpha)} \\
&\geq \frac{\log\left(\frac{1}{n} \sum_{i=1}^n \exp(-q_i(\boldsymbol{\theta}))\right)}{\|\mathbf{h}\|_2} + \mathcal{O}(\alpha^{-cL} \log A) \\
&\geq \frac{\log \frac{1}{\mathcal{L}(\boldsymbol{\theta})}}{\|\mathbf{h}\|_2} + \mathcal{O}(\alpha^{-cL} \log A) \rightarrow \gamma_{\text{ntk}}.
\end{aligned}$$

Therefore, we have $\min_{i \in [n]} \langle \nabla q_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \frac{\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})}{\|\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})\|_2} \rangle \rightarrow \gamma_{\text{ntk}}$. By the uniqueness of the max-margin solution of [\(K1\)](#), it must hold that $\frac{\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})}{\|\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})\|_2} \rightarrow \mathbf{h}_{\text{ntk}}^*$ as $\alpha \rightarrow +\infty$.

Then, we can approximate the norm of \mathbf{h} as

$$\|\mathbf{h}\|_2 = \frac{\log \frac{1}{\mathcal{L}(\boldsymbol{\theta})}}{\gamma_{\text{ntk}} + o(1)} = \frac{\log \frac{\alpha^c}{\lambda} + O(\log \log A)}{\gamma_{\text{ntk}} + o(1)} = \frac{1}{\gamma_{\text{ntk}}} \log \frac{\alpha^c}{\lambda} (1 + o(1)).$$

So \mathbf{h} can be approximated by

$$\mathbf{h} = \|\mathbf{h}\|_2 \frac{\mathbf{h}}{\|\mathbf{h}\|_2} = (1 + o(1)) \frac{1}{\gamma_{\text{ntk}}} \log \frac{\alpha^c}{\lambda} \mathbf{h}_{\text{ntk}}^*,$$

which completes the proof. \square

Now we are ready to prove [Theorem 3.4](#).

Proof: For any input \mathbf{x} , by a simple extension of [Lemma C.6](#) to $f(\cdot; \mathbf{x})$ we have

$$f(\boldsymbol{\theta}; \mathbf{x}) = \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h} \rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \|\mathbf{h}\|_2.$$

Normalizing by $Z(\alpha)$ and taking limits, we have

$$\frac{f(\boldsymbol{\theta}; \mathbf{x})}{Z(\alpha)} = \left\langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \frac{\mathbf{h}}{Z(\alpha)} \right\rangle + \mathcal{O}(\alpha^{-L} \log A) \cdot e^{L\lambda t} \cdot \frac{\|\mathbf{h}\|_2}{Z(\alpha)} \rightarrow \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h}_{\text{ntk}}^* \rangle,$$

which follows from [Theorem C.9](#). \square

C.3 RICH REGIME

Now we proceed to prove [Theorem 3.5](#) for the rich regime. First, we derive the following bound for the training loss at the end of the kernel regime.

Lemma C.10 *At $t = \frac{1}{\lambda}(\log \alpha - \frac{1}{L} \log \log A + \Delta T)$, we have $\mathcal{L}_\lambda(\boldsymbol{\theta}(t)) = \mathcal{O}(\lambda(\log \frac{1}{\lambda})^{2/L})$.*

Proof: By [Lemma C.4](#), we can easily obtain the bound

$$\mathcal{L}(\boldsymbol{\theta}(t)) = \mathcal{O}\left(\frac{\log A}{A} \exp(2(L-1)(\log \alpha - \frac{1}{L} \log \log A))\right) = \mathcal{O}(\lambda(\log A)^{-1+2/L}).$$

Also we have the norm bound $\|\boldsymbol{\theta}(t)\|_2 = \mathcal{O}(\alpha e^{-\lambda t}) = \mathcal{O}(\log A)^{1/L}$ since it is in the kernel regime. We can conclude the proof by noting that $\mathcal{L}_\lambda = \mathcal{L} + \|\boldsymbol{\theta}\|_2^2$ and $\log A = \mathcal{O}(\log \frac{1}{\lambda})$. \square

The loss bound then implies the following norm bound.

Lemma C.11 *If $\mathcal{L}_\lambda(\boldsymbol{\theta}(0)) = \mathcal{O}(\lambda(\log \frac{1}{\lambda})^{2/L})$, then for all $t \geq 0$,*

$$\|\boldsymbol{\theta}(t)\|_2 = \Theta((\log \frac{1}{\lambda})^{1/L}).$$

Proof: For all $t \geq 0$, the following relationship can be deduced from the monotonicity of \mathcal{L}_λ , $r_i(\boldsymbol{\theta}(t)) \leq n\mathcal{L}_\lambda(\boldsymbol{\theta}(t))$, and the Lipschitzness of f_i :

$$\Omega(\log(1/\lambda)) \leq \log \frac{1}{n\mathcal{L}_\lambda(\boldsymbol{\theta}(0))} \leq \log \frac{1}{n\mathcal{L}_\lambda(\boldsymbol{\theta}(t))} \leq q_i(\boldsymbol{\theta}(t)) \leq B_0 \|\boldsymbol{\theta}(t)\|_2^L.$$

So $\|\boldsymbol{\theta}(t)\|_2 = \Omega((\log \frac{1}{\lambda})^{1/L})$. Also note $\|\boldsymbol{\theta}(t)\|_2^2 \leq \frac{2}{\lambda} \mathcal{L}_\lambda(\boldsymbol{\theta}(t)) \leq \frac{2}{\lambda} \mathcal{L}_\lambda(\boldsymbol{\theta}(0)) = \mathcal{O}((\log \frac{1}{\lambda})^{1/L})$. So we have the tight norm bound $\|\boldsymbol{\theta}(t)\|_2 = \Theta((\log \frac{1}{\lambda})^{1/L})$. \square

To prove [Theorem 3.5](#), it suffices to prove the following.

Theorem C.12 *For any starting point $\boldsymbol{\theta}_0(\alpha)$ satisfying $\mathcal{L}_\lambda(\boldsymbol{\theta}_0(\alpha)) = \mathcal{O}(\lambda(\log \frac{1}{\lambda})^{2/L})$ and a time function $T(\alpha) = \Omega(\frac{1}{\lambda} \log \frac{1}{\lambda})$, given any sequence $\{\alpha_k\}$ with $\alpha_k \rightarrow +\infty$, there exists a sequence $\{t_k\}_{k \geq 1}$ such that $t_k \leq T(\alpha)$, and every limit point of $\{\frac{\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))}{\|\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))\|_2} : k \geq 0\}$ is along the direction of a KKT point of [\(R1\)](#).*

Proof: For all $\alpha > 0$, $0 \leq \mathcal{L}_\lambda(\boldsymbol{\theta}(T(\alpha))) = \mathcal{L}_\lambda(\boldsymbol{\theta}_0) - \int_0^{T(\alpha)} \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}(t))\|_2^2 dt$. Since $\mathcal{L}_\lambda(\boldsymbol{\theta}_0(\alpha)) = \mathcal{O}(\lambda(\log A)^{2/L})$, we have the bound $\int_0^{T(\alpha)} \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}(t))\|_2^2 dt \leq \mathcal{O}(\lambda(\log \frac{1}{\lambda})^{2/L})$, which implies that there exists a time τ_α such that $\|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}(\tau_\alpha))\|_2^2 \leq \mathcal{O}(\lambda^2(\log \frac{1}{\lambda})^{-(1-2/L)})$.

Now let $t_k = \tau_{\alpha_k}$. It suffices to prove in the case where $\frac{\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))}{\|\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))\|_2} \in \mathcal{S}^{d-1}$ converges; otherwise we can do the same prove for any convergent subsequence. Denote $\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))$ as $\boldsymbol{\theta}^{(k)}$ and $\frac{\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))}{\|\boldsymbol{\theta}(t_k; \boldsymbol{\theta}_0(\alpha_k))\|_2}$ as $\bar{\boldsymbol{\theta}}^{(k)}$ for short. We claim that the limit $\bar{\boldsymbol{\theta}} := \lim_{k \rightarrow +\infty} \bar{\boldsymbol{\theta}}^{(k)}$ is along a KKT-margin direction of [\(R1\)](#).

First, according to our choice of t_k , we have

$$\left\| \frac{1}{n} \sum_{i=1}^n r_i(\boldsymbol{\theta}^{(k)}) \nabla q_i(\boldsymbol{\theta}^{(k)}) - \lambda \boldsymbol{\theta}^{(k)} \right\|_2 \leq \mathcal{O}(\lambda(\log A)^{-(1/2-1/L)}). \quad (21)$$

By [Lemma C.11](#), $\lambda \|\boldsymbol{\theta}^{(k)}\|_2 = \mathcal{O}(\lambda(\log \frac{1}{\lambda})^{1/L})$. Now we divide $\lambda \|\boldsymbol{\theta}^{(k)}\|_2$ on both sides of (21):

$$\left\| \sum_{i=1}^n \frac{\|\boldsymbol{\theta}\|_2^{L-2}}{n\lambda} r_i(\boldsymbol{\theta}^{(k)}) \nabla q_i(\bar{\boldsymbol{\theta}}^{(k)}) - \bar{\boldsymbol{\theta}}^{(k)} \right\|_2 \leq \mathcal{O}\left(\frac{(\log A)^{-(1/2-1/L)}}{(\log \frac{1}{\lambda})^{1/L}}\right). \quad (22)$$

Let $\boldsymbol{\mu}_k$ be a vector with coordinates $\mu_{k,i} := \frac{\|\boldsymbol{\theta}\|_2^{L-2}}{n\lambda} r_i(\boldsymbol{\theta}^{(k)})$. We can show that $\mu_{k,i} = \mathcal{O}(1)$ by noticing $f_i(\bar{\boldsymbol{\theta}}^{(k)}) = \frac{1}{\|\boldsymbol{\theta}\|_2^L} f_i(\boldsymbol{\theta}^{(k)}) = \Omega(1)$ and

$$\begin{aligned} \sum_{i=1}^n \mu_{k,i} y_i f_i(\bar{\boldsymbol{\theta}}^{(k)}) &\leq 1 + \left| \sum_{i=1}^n \mu_{k,i} y_i f_i(\bar{\boldsymbol{\theta}}^{(k)}) - 1 \right| \\ &= 1 + \left| \left\langle \sum_{i=1}^n \mu_{k,i} y_i \nabla f_i(\bar{\boldsymbol{\theta}}^{(k)}) - \bar{\boldsymbol{\theta}}^{(k)}, \bar{\boldsymbol{\theta}}^{(k)} \right\rangle \right| \\ &= 1 + \mathcal{O}\left(\frac{1}{\log \frac{1}{\lambda}} (\log A)^{-(1/2-1/L)}\right) = \mathcal{O}(1). \end{aligned}$$

Then, let $\{(\alpha_{k_p}, t_{k_p})\}_{p \geq 1}$ be a subsequence of $\{(\alpha_k, t_k)\}_{k \geq 1}$ so that $\boldsymbol{\mu}_{k_p} \in \mathbb{R}^n$ converges. Let $\bar{\boldsymbol{\mu}} := \lim_{p \rightarrow +\infty} \boldsymbol{\mu}_{k_p}$ be the corresponding limit. We can take limit $k_p \rightarrow +\infty$ on both sides of (22) and obtain $\sum_{i=1}^n \bar{\mu}_i y_i \nabla f_i(\bar{\boldsymbol{\theta}}) = \bar{\boldsymbol{\theta}}$.

Let $i_* \in [n]$ be an index such that $q_{i_*}(\bar{\boldsymbol{\theta}}) = q_{\min}(\bar{\boldsymbol{\theta}})$. We can verify that $\bar{\mu}_i = 0$ for all $i \in [n]$ with $f_i(\bar{\boldsymbol{\theta}}) > q_{\min}(\bar{\boldsymbol{\theta}})$ by noting $\frac{\mu_{k,i}}{\mu_{k,i_*}} = \exp(q_{i_*}(\boldsymbol{\theta}^{(k)}) - f_i(\boldsymbol{\theta}^{(k)})) = \exp(\|\boldsymbol{\theta}\|_2^L (q_{i_*}(\boldsymbol{\theta}) - q_i(\bar{\boldsymbol{\theta}}))) \rightarrow 0$ and $\bar{\boldsymbol{\mu}} \neq \mathbf{0}$. \square

Finally, [Theorem 3.5](#) can be proved by simply combining [Lemma C.10](#) with [Theorem C.12](#).

D PROOFS FOR REGRESSION WITH L2 REGULARIZATION

The analyses for the regression task are similar to those for the classification task, but the use of squared loss simplifies the analysis.

We define $s_i(\boldsymbol{\theta}) := f_i(\boldsymbol{\theta}) - y_i$. In the regression setting, the gradients of the vanilla and regularized square loss can be written as

$$\nabla \mathcal{L}(\boldsymbol{\theta}) = 2 \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}), \quad (23)$$

and

$$\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}) = 2 \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}) + \lambda \boldsymbol{\theta}. \quad (24)$$

D.1 KERNEL REGIME

Now we present the proof of [Theorem 3.9](#) for the kernel regime. Following [Section 3.3.1](#), we define $\mathbf{h}_{\text{ntk}}^*$ as the solution to [\(K2\)](#). Additionally, let $\Phi(\boldsymbol{\theta}) \in \mathbb{R}^{D \times n}$ be the matrix where the i -th column is $\nabla f_i(\boldsymbol{\theta})$. Let $\mathbf{K}(\boldsymbol{\theta}) := \Phi(\boldsymbol{\theta})^\top \Phi(\boldsymbol{\theta})$ be the NTK matrix at $\boldsymbol{\theta}$. Let $\nu_{\text{ntk}} > 0$ be the minimum eigenvalue of $\mathbf{K}(\bar{\boldsymbol{\theta}}_{\text{init}})$. Let ϵ_{max} be a small constant so that $\lambda_{\min}(\mathbf{K}(\boldsymbol{\theta})) > \frac{\nu_{\text{ntk}}}{2}$ holds for all $\|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 < \epsilon_{\text{max}}$. Let $T_{\text{max}} := \inf \left\{ t \geq 0 : \left\| \frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 > \epsilon_{\text{max}} \right\}$, which is the time that $\boldsymbol{\theta}(t)$ moves far enough to change the NTK features significantly.

First, we derive an upper bound for the loss convergence.

Lemma D.1 For all $0 \leq t \leq T_{\text{max}}$,

$$\frac{d\mathcal{L}}{dt} \leq -\frac{\nu_{\text{ntk}}}{2n} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \mathcal{L} + \frac{\lambda L}{\sqrt{n}} \|\mathbf{y}\|_2 \sqrt{\mathcal{L}}, \quad (25)$$

$$\mathcal{L}(\boldsymbol{\theta}(t)) \leq \max \left\{ \frac{1}{n} \|\mathbf{y}\|_2^2 \exp \left(-\frac{\nu_{\text{ntk}} A}{8n(L-1)} \left(1 - e^{-2(L-1)\lambda t} \right) \right), \frac{16nL^2 \|\mathbf{y}\|_2^2 \lambda^2}{\nu_{\text{ntk}}^2 \alpha^{4(L-1)}} e^{4(L-1)\lambda t} \right\}. \quad (26)$$

Proof: By the update rule of $\boldsymbol{\theta}$, the following holds for all $t > 0$:

$$\frac{d\mathcal{L}}{dt} = \left\langle \nabla \mathcal{L}(\boldsymbol{\theta}), \frac{d\boldsymbol{\theta}}{dt} \right\rangle = - \underbrace{\left\| \frac{1}{n} \sum_{i=1}^n s_i \nabla f_i \right\|_2^2}_{=: V_1} - \underbrace{\frac{\lambda L}{n} \sum_{i=1}^n s_i f_i}_{=: V_2}.$$

For V_1 , we have

$$\begin{aligned} -V_1 &= \left\| \frac{1}{n} \sum_{i=1}^n r_i \nabla f_i(\boldsymbol{\theta}) \right\|_2^2 = \frac{1}{n^2} \mathbf{s}(\boldsymbol{\theta})^\top \mathbf{K}(\boldsymbol{\theta}) \mathbf{s}(\boldsymbol{\theta}) \\ &= \frac{1}{n^2} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \mathbf{s}(\boldsymbol{\theta})^\top \mathbf{K} \left(\frac{e^{\lambda t}}{\alpha} \boldsymbol{\theta} \right) \mathbf{s}(\boldsymbol{\theta}) \\ &\geq \frac{\nu_{\text{ntk}}}{2n^2} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \|\mathbf{s}(\boldsymbol{\theta})\|_2^2 \\ &= \frac{\nu_{\text{ntk}}}{2n} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} \mathcal{L}. \end{aligned}$$

For V_2 , we have $\frac{1}{n} \sum_{i=1}^n s_i f_i = \frac{1}{n} \sum_{i=1}^n (s_i^2 + s_i y_i) \geq \frac{1}{n} \sum_{i=1}^n s_i y_i \geq -\frac{1}{\sqrt{n}} \|\mathbf{y}\|_2 \sqrt{\mathcal{L}}$. So $V_2 \leq \frac{\lambda L}{\sqrt{n}} \|\mathbf{y}\|_2 \sqrt{\mathcal{L}}$. Combining the inequalities for V_1 and V_2 together proves [\(25\)](#).

For proving [\(26\)](#), we first divide by \mathcal{L} on both sides of [\(25\)](#) to get

$$\frac{d}{dt} \log \mathcal{L} \leq -\frac{\nu_{\text{ntk}}}{2n} \alpha^{2(L-1)} e^{-2(L-1)\lambda t} + \frac{\lambda L}{\sqrt{n}} \|\mathbf{y}\|_2 \cdot \frac{1}{\sqrt{\mathcal{L}}}.$$

Let $\mathcal{E} := \{t \in [0, T_{\max}] : \frac{\lambda L}{\sqrt{n}} \|\mathbf{y}\|_2 \cdot \frac{1}{\sqrt{\mathcal{L}}} \geq \frac{\nu_{\text{ntk}}}{4n} \alpha^{2(L-1)} e^{-2(L-1)\lambda t}\}$. For all $t \in [0, T_{\max}]$, if $t \in \mathcal{E}$, then

$$\begin{aligned} \mathcal{L} &\leq \frac{\lambda^2 L^2}{n} \|\mathbf{y}\|_2^2 \cdot \frac{1}{\left(\frac{\nu_{\text{ntk}}}{4n} \alpha^{2(L-1)} e^{-2(L-1)\lambda t}\right)^2} \\ &= \frac{16nL^2 \|\mathbf{y}\|_2^2}{\nu_{\text{ntk}}^2} \cdot \frac{\lambda^2}{\alpha^{4(L-1)}} e^{4(L-1)\lambda t}. \end{aligned}$$

Otherwise, let t' be the largest number in $\mathcal{E} \cup \{0\}$ that is smaller than t . Then

$$\begin{aligned} \log \mathcal{L}(\boldsymbol{\theta}(t)) &= \log \mathcal{L}(\boldsymbol{\theta}(t')) + \int_{t'}^t \frac{d}{d\tau} \log \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau \\ &\leq \log \mathcal{L}(\boldsymbol{\theta}(t')) - \int_{t'}^t \left(\frac{\nu_{\text{ntk}}}{4n} \alpha^{2(L-1)} e^{-2(L-1)\lambda \tau}\right) d\tau \\ &\leq \log \mathcal{L}(\boldsymbol{\theta}(t')). \end{aligned}$$

If $t' \in \mathcal{E}$, then $\mathcal{L}(\boldsymbol{\theta}(t)) \leq \mathcal{L}(\boldsymbol{\theta}(t')) \leq \frac{16nL^2 \|\mathbf{y}\|_2^2}{\nu_{\text{ntk}}^2} \cdot \frac{\lambda^2}{\alpha^{4(L-1)}} e^{4(L-1)\lambda t'} \leq \frac{16nL^2 \|\mathbf{y}\|_2^2}{\nu_{\text{ntk}}^2} \cdot \frac{\lambda^2}{\alpha^{4(L-1)}} e^{4(L-1)\lambda t}$. Otherwise, $t' = 0$ and we have

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}(t)) &\leq \mathcal{L}(\boldsymbol{\theta}(0)) \exp\left(-\int_0^t \left(\frac{\nu_{\text{ntk}}}{4n} \alpha^{2(L-1)} e^{-2(L-1)\lambda \tau}\right) d\tau\right) \\ &= \frac{1}{n} \|\mathbf{y}\|_2^2 \exp\left(-\frac{\nu_{\text{ntk}} A}{8n(L-1)} \left(1 - e^{-2(L-1)\lambda t}\right)\right) \end{aligned}$$

which concludes the proof. \square

Now we derive a lower bound for the time that the dynamics stay in the kernel regime.

Lemma D.2 *There exists a constant ΔT such that $T_{\max} \geq \frac{1}{\lambda} (\log \alpha + \Delta T)$ when α is sufficiently large. Furthermore, for all $t \leq \frac{1}{\lambda} (\log \alpha + \Delta T)$,*

$$\left\| \frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 = \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t}.$$

Proof: By definition of T_{\max} , to prove the first claim, it suffices to show that there exists a constant ΔT such that $\left\| \frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 \leq \epsilon_{\max}$ for all $t \leq \min\{\frac{1}{\lambda} (\log \alpha + \Delta T), T_{\max}\}$.

When $t \leq T_{\max}$, we have the following gradient upper bound:

$$\begin{aligned} \|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_2 &= \left\| \frac{1}{n} \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}(t)) \right\|_2 \\ &\leq \frac{1}{n} \|\mathbf{s}(\boldsymbol{\theta})\|_2 \|\Phi(\boldsymbol{\theta})\|_2 \\ &= \frac{1}{\sqrt{n}} \alpha^{L-1} e^{-(L-1)\lambda t} \|\Phi(\frac{1}{\alpha} e^{-\lambda t} \boldsymbol{\theta})\|_2 \cdot \sqrt{\frac{1}{n} \|\mathbf{s}(\boldsymbol{\theta})\|_2^2} \\ &\leq G_{\text{ntk}} \alpha^{L-1} e^{-(L-1)\lambda t} \sqrt{\mathcal{L}}. \end{aligned}$$

By update rule, $e^{\lambda t} \boldsymbol{\theta}(t) - \alpha \bar{\boldsymbol{\theta}}_{\text{init}} = \int_0^t e^{\lambda \tau} \nabla \mathcal{L}(\boldsymbol{\theta}(\tau)) d\tau$. Applying the gradient upper bound gives

$$\begin{aligned} \frac{1}{\alpha} \left\| e^{\lambda t} \boldsymbol{\theta}(t) - \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right\|_2 &\leq \frac{1}{\alpha} \int_0^t \left\| e^{\lambda \tau} \nabla \mathcal{L}(\boldsymbol{\theta}(\tau)) \right\|_2 d\tau \\ &\leq \frac{1}{\alpha} \int_0^t e^{\lambda \tau} G_{\text{ntk}} \alpha^{L-1} e^{-(L-1)\lambda \tau} \sqrt{\mathcal{L}(\boldsymbol{\theta}(\tau))} d\tau \\ &= G_{\text{ntk}} \alpha^{L-2} \int_0^t e^{-(L-2)\lambda \tau} \sqrt{\mathcal{L}(\boldsymbol{\theta}(\tau))} d\tau. \end{aligned}$$

Now we substitute \mathcal{L} in $\int_0^t e^{-(L-2)\lambda\tau} \sqrt{\mathcal{L}(\boldsymbol{\theta}(\tau))} d\tau$ with the loss upper bound in [Lemma D.1](#). Then for all $t \leq T_{\max}$, we have $\int_0^t e^{-(L-2)\lambda\tau} \sqrt{\mathcal{L}(\boldsymbol{\theta}(\tau))} d\tau \leq \max\{I_1, I_2\}$, where

$$I_1 := \int_0^t \frac{1}{\sqrt{n}} \|\mathbf{y}\|_2 \exp\left(-\frac{\nu_{\text{ntk}} A}{16n(L-1)} \left(1 - e^{-2(L-1)\lambda\tau}\right) - (L-2)\lambda\tau\right),$$

$$I_2 := \int_0^t \frac{4\sqrt{n}L\|\mathbf{y}\|_2}{\nu_{\text{ntk}}} \cdot \frac{\lambda}{\alpha^{2(L-1)}} e^{L\lambda\tau} d\tau.$$

We now bound I_1 and I_2 . For I_1 , we divide the integral into two parts. Let $t_0 = \frac{1}{20(L-1)\lambda}$, for $\tau \in [0, t_0]$, we have $1 - e^{-2(L-1)\lambda\tau} \geq \frac{1}{2} \cdot 2(L-1)\lambda\tau = (L-1)\lambda\tau$, so that

$$\int_0^{t_0} \exp\left(-\frac{\nu_{\text{ntk}} A}{16n(L-1)} \left(1 - e^{-2(L-1)\lambda\tau}\right) - (L-2)\lambda\tau\right)$$

$$\leq \int_0^{t_0} \exp\left(-\frac{\nu_{\text{ntk}} A \lambda \tau}{16n}\right) d\tau = \mathcal{O}\left(\frac{1}{A\lambda}\right) = \mathcal{O}\left(\frac{1}{\alpha^{2(L-1)}}\right).$$

On the other hand, we have

$$\int_{t_0}^t \exp\left(-\frac{\nu_{\text{ntk}} A}{16n(L-1)} \left(1 - e^{-2(L-1)\lambda\tau}\right) - (L-2)\lambda\tau\right)$$

$$\leq \exp(-\Omega(A)) \cdot \int_{t_0}^t \exp(-(L-2)\lambda\tau) d\tau \leq \exp(-\Omega(A)) \cdot \mathcal{O}(\lambda^{-1}) = \mathcal{O}\left(\frac{1}{\alpha^{2(L-1)}}\right).$$

Thus, we have

$$I_1 = \mathcal{O}\left(\frac{1}{\alpha^{2(L-1)}}\right).$$

For I_2 , it holds for all $t \leq T_{\max}$ that

$$I_2 = \mathcal{O}\left(\frac{1}{\alpha^{2(L-1)}}\right) \cdot e^{L\lambda t}.$$

Putting all these together, we have

$$\frac{1}{\alpha} \|e^{\lambda t} \boldsymbol{\theta}(t) - \alpha \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 = \mathcal{O}(\alpha^{L-2}) \cdot \mathcal{O}\left(\frac{1}{\alpha^{2(L-1)}}\right) \cdot e^{L\lambda t}$$

$$= \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t}. \quad (27)$$

Therefore, there exists a constant $C > 0$ such that when α is sufficiently large, $\left\|\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}}\right\|_2 \leq C \alpha^{-L} e^{L\lambda t}$ for all $t \leq \min\{T_{\max}, \mathcal{O}(\frac{1}{\lambda} \log \alpha)\}$. Setting $\Delta T := \frac{1}{L} \log \frac{\epsilon_{\max}}{C}$, we obtain the following bound for all $t \leq \min\{T_{\max}, \frac{1}{\lambda}(\log \alpha + \Delta T)\}$:

$$\left\|\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}}\right\|_2 \leq C \alpha^{-L} \exp\left(L \log \alpha + \log \frac{\epsilon_{\max}}{C}\right) = \epsilon_{\max},$$

which implies $T_{\max} \geq \frac{1}{\lambda}(\log \alpha + \Delta T)$. Combining this with [\(13\)](#) prove the second claim. \square

Let $\boldsymbol{\delta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) := \left(\frac{1}{\alpha} e^{\lambda t} \boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) - \bar{\boldsymbol{\theta}}_{\text{init}}\right)$ and $\mathbf{h}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) := \alpha^L e^{-L\lambda t} \boldsymbol{\delta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})$. In the following, we show via a series of lemmas that for any \mathbf{v} that is in the orthogonal complement of Φ , $\frac{\langle \boldsymbol{\delta}, \mathbf{v} \rangle}{\|\boldsymbol{\delta}\|_2} \rightarrow 0$, which is a crucial property for deriving the implicit bias.

Lemma D.3 *For any vector \mathbf{v} that lies in the orthogonal space of the column space of Φ , we have $\langle \boldsymbol{\delta}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{v} \rangle = \mathcal{O}(\alpha^{-2cL})$ where $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$.*

Proof: By update rule,

$$\frac{d\boldsymbol{\delta}}{dt} = \frac{1}{\alpha} e^{\lambda t} \left(-2 \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})) \right).$$

Then we have

$$\begin{aligned} \frac{d\langle \boldsymbol{\delta}, \mathbf{v} \rangle}{dt} &= -2\frac{1}{\alpha}e^{\lambda t} \left\langle \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})), \mathbf{v} \right\rangle \\ &= -2\frac{1}{\alpha}e^{\lambda t} \left\langle \sum_{i=1}^n s_i (\nabla f_i(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})) - \nabla f_i(\alpha e^{-\lambda t} \bar{\boldsymbol{\theta}}_{\text{init}})), \mathbf{v} \right\rangle \\ &= -2\alpha^{L-2}e^{-(L-2)\lambda t} \left\langle \sum_{i=1}^n s_i (\nabla f_i(\alpha^{-1}e^{\lambda t}\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})) - \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}})), \mathbf{v} \right\rangle \end{aligned}$$

where the second equation is due to our choice of \mathbf{v} and the last equation follows from the homogeneity of f_i . Now, since f_i is locally smooth in a neighbourhood of $\bar{\boldsymbol{\theta}}_{\text{init}}$ and $\|\frac{1}{\alpha}e^{\lambda t}\boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 = \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t}$ by [Lemma D.2](#), we have

$$\|\nabla f_i(\alpha^{-1}e^{\lambda t}\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) - \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \mathcal{O}(\alpha^{-L}e^{L\lambda t}).$$

Let $t_0 = \frac{1}{20(L-1)\lambda}$, then for $t \in [0, t_0]$, we have $1 - e^{2(L-1)\lambda\tau} \geq (L-1)\lambda\tau$, so that

$$\begin{aligned} \frac{d\langle \boldsymbol{\delta}, \mathbf{v} \rangle}{dt} &= \mathcal{O}\left(\alpha^{-2}e^{2\lambda t} \sqrt{\mathcal{L}(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}))}\right) \\ &= \mathcal{O}\left(\alpha^{-2}e^{2\lambda t} \max\left\{\exp\left(-\frac{\nu_{\text{ntk}}A\lambda t}{16n} + \lambda t\right), \alpha^{-2(L-1)}\lambda e^{2(L-1)\lambda t}\right\}\right) \\ &= \mathcal{O}\left(\max\left\{\alpha^{-2} \exp\left(-\frac{\nu_{\text{ntk}}A\lambda t}{32n}\right), \lambda\alpha^{-2L}e^{2L\lambda t}\right\}\right) \end{aligned}$$

for sufficiently large α , since $A = \frac{\alpha^{2(L-1)}}{\lambda}$. On the other hand, for $t \in [t_0, T_c^-(\alpha)]$, the second term in (26) dominates the first term, so that

$$\begin{aligned} \frac{d\langle \boldsymbol{\delta}, \mathbf{v} \rangle}{dt} &= \mathcal{O}\left(\alpha^{-2}e^{2\lambda t} \sqrt{\mathcal{L}(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}))}\right) \\ &= \mathcal{O}\left(\alpha^{-2}e^{2\lambda t} \cdot \alpha^{-2(L-1)}\lambda e^{2(L-1)\lambda t}\right) = \mathcal{O}(\lambda\alpha^{-2L}e^{2L\lambda t}). \end{aligned}$$

Hence,

$$\begin{aligned} \langle \boldsymbol{\delta}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{v} \rangle &= \mathcal{O}\left(\max\left\{\alpha^{-2} \int_0^{T_c^-(\alpha)} \exp\left(-\frac{\nu_{\text{ntk}}A\lambda t}{32n}\right) dt, \lambda\alpha^{-2L} \int_0^{T_c^-(\alpha)} e^{2L\lambda t} dt\right\}\right) \\ &= \mathcal{O}\left(\max\left\{\alpha^{-2} \frac{1}{A\lambda}, \alpha^{-2L}\alpha^{2L(1-c)}\right\}\right) = \mathcal{O}(\alpha^{-2cL}), \end{aligned}$$

which proves the claim. \square

Lemma D.4 For $\boldsymbol{\theta} = \boldsymbol{\theta}(t)$ with $t \leq \frac{1}{\lambda}(\log \alpha + \Delta T)$ as defined in [Lemma D.2](#), we have

$$f\left(\frac{1}{\alpha}e^{\lambda t}\boldsymbol{\theta}; \mathbf{x}\right) = \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \boldsymbol{\delta} \rangle + \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t} \|\boldsymbol{\delta}\|_2, \quad (28)$$

$$f(\boldsymbol{\theta}; \mathbf{x}) = \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h} \rangle + \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t} \|\mathbf{h}\|_2. \quad (29)$$

Proof: By [Lemma D.2](#) and Taylor expansion, we have

$$\begin{aligned} f\left(\frac{1}{\alpha}e^{\lambda t}\boldsymbol{\theta}; \mathbf{x}\right) &= f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}) + \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \boldsymbol{\delta} \rangle + \mathcal{O}(\|\boldsymbol{\delta}\|_2^2) \\ &= f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}) + \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \boldsymbol{\delta} \rangle + \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t} \|\boldsymbol{\delta}\|_2 \\ &= \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \boldsymbol{\delta} \rangle + \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda t} \|\boldsymbol{\delta}\|_2, \end{aligned}$$

which proves (28). Combining this with the L -homogeneity of q_i proves (29). \square

Lemma D.5 $\|\mathbf{h}(T_c^-(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \Theta(1)$.

Proof: By Lemma D.2 we have $\|\delta(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \mathcal{O}(e^{-cL})$, so that $\|\mathbf{h}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \alpha^L e^{-L\lambda T_c^-} \|\delta(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \mathcal{O}(1)$.

On the other hand, recall that $T_c^-(\alpha) = \frac{1-c}{\lambda} \log \alpha$, so by (26) we have

$$|s_i| \leq \sqrt{n\mathcal{L}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}}))} = \mathcal{O}\left(\max\left\{e^{-\Omega(A)}, \lambda\alpha^{-2(L-1)c}\right\}\right) = \mathcal{O}\left(\alpha^{-2(L-1)c}\right).$$

Thus, for sufficiently large α , we have $|f_i(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}}))| \geq \frac{1}{2}y_i$.

By Lemma D.4, we have that

$$\frac{1}{2}y_i \leq \|\nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 \|\mathbf{h}\|_2 + \mathcal{O}(\alpha^{-cL}) \|\mathbf{h}\|_2 \Rightarrow \|\mathbf{h}\|_2 = \Omega(1)$$

as desired. \square

Corollary D.6 $\lim_{\alpha \rightarrow 0} \left\langle \frac{\delta}{\|\delta\|_2}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{v} \right\rangle = 0$.

Proof: By the previous lemma, we have $\|\delta(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \Omega\left(\alpha^{-L} e^{L\lambda T_c^-(\alpha)}\right) = \Omega(\alpha^{-cL})$.

On the other hand, $\langle \delta(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{v} \rangle = \mathcal{O}(\alpha^{-2cL})$ by Lemma D.3. The conclusion immediately follows. \square

Theorem D.7 For any constant $c \in (0, 1)$, letting $T_c^-(\alpha) := \frac{1-c}{\lambda} \log \alpha$, it holds that

$$\forall \mathbf{x} \in \mathbb{R}^d: \quad \lim_{\alpha \rightarrow +\infty} f(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}}); \mathbf{x}) = \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h}_{\text{ntk}}^* \rangle.$$

Proof: By Lemmas D.4 and D.5 we have

$$\begin{aligned} & f(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}}); \mathbf{x}) \\ &= \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) \rangle + \mathcal{O}(\alpha^{-L}) \cdot e^{L\lambda T_c^-(\alpha)} \|\mathbf{h}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 \quad (30) \\ &= \langle \nabla f(\bar{\boldsymbol{\theta}}_{\text{init}}; \mathbf{x}), \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) \rangle + \mathcal{O}(\alpha^{-cL}) \end{aligned}$$

By Lemma D.1 we have for $\forall i \in [n]$, $|f_i(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) - y_i| = \mathcal{O}(\alpha^{-2(L-1)})$. Since $y_i = \langle \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}_{\text{ntk}}^* \rangle$, we have

$$\lim_{\alpha \rightarrow +\infty} \langle \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) - \mathbf{h}_{\text{ntk}}^* \rangle = 0. \quad (31)$$

By Corollary D.6 and Lemma D.5 we have that for any direction \mathbf{v} orthogonal to the column space of Φ , we have $\lim_{\alpha \rightarrow +\infty} \langle \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})), \mathbf{v} \rangle = 0$. Let \mathbf{P}_Φ be the projection operator onto the column space of Φ , then

$$\lim_{\alpha \rightarrow +\infty} (\mathbf{I} - \mathbf{P}_\Phi) \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) = 0. \quad (32)$$

Combined with (31), we deduce that

$$\lim_{\alpha \rightarrow +\infty} \langle \nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), \mathbf{P}_\Phi \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) - \mathbf{h}_{\text{ntk}}^* \rangle = 0.$$

The above holds for all $\nabla f_i(\bar{\boldsymbol{\theta}}_{\text{init}}), i \in [d]$, which are exactly the columns of Φ . Note also that $\mathbf{h}_{\text{ntk}}^*$ also lies in the column space of Φ , so we actually have $\lim_{\alpha \rightarrow +\infty} \|\mathbf{P}_\Phi \mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) - \mathbf{h}_{\text{ntk}}^*\|_2 = 0$. Hence, we can use (32) to deduce that $\lim_{\alpha \rightarrow +\infty} \|\mathbf{h}(\boldsymbol{\theta}(T_c^-(\alpha); \alpha\bar{\boldsymbol{\theta}}_{\text{init}})) - \mathbf{h}_{\text{ntk}}^*\|_2 = 0$. Finally, plugging this into (30) gives the desired result. \square

D.2 RICH REGIME

Now we proceed to prove Theorem 3.10 for the rich regime. First, we derive a norm bound.

Lemma D.8 For any constant $c > 0$, let $T_c^+(\alpha) := \frac{1+c}{\lambda} \log \alpha$, then we have that

$$\max_{\frac{1}{\lambda} \log \alpha \leq t \leq T_c^+(\alpha)} \|\boldsymbol{\theta}(t; \alpha\bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \mathcal{O}(1).$$

Proof: We know from [Lemmas D.1](#) and [D.2](#) that

$$\mathcal{L} \left(\boldsymbol{\theta} \left(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right) \right) \leq \max \left\{ \frac{1}{n} \|y\|_2^2 \exp \left(-\frac{\nu_{\text{ntk}} A}{16n(L-1)} \right), \frac{16nL^2 \|y\|_2^2 \lambda^2}{\nu_{\text{ntk}}^2} \right\} \leq \mathcal{O}(\lambda^2).$$

Moreover, by [Lemma D.2](#) we know that $\|\boldsymbol{\theta} \left(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right)\|_2 = \mathcal{O}(1)$, so that

$$\mathcal{L}_\lambda \left(\boldsymbol{\theta} \left(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right) \right) = \mathcal{L} \left(\boldsymbol{\theta} \left(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right) \right) + \frac{\lambda}{2} \left\| \boldsymbol{\theta} \left(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right) \right\|_2^2 = \mathcal{O}(\lambda).$$

Since $\{\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})\}_{t \geq 0}$ is the trajectory of GF on \mathcal{L}_λ , we know that for $\forall t \geq \frac{1}{\lambda} \log \alpha$, we have $\mathcal{L}_\lambda(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})) = \mathcal{O}(\lambda)$ as well. This implies that $\|\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \mathcal{O}(1)$ as desired. \square

Now we prove [Theorem 3.10](#).

Proof: By [Lemma C.5](#), we know that for $t = \frac{1}{\lambda} \log \alpha$, we have $\|\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}) - \bar{\boldsymbol{\theta}}_{\text{init}}\|_2 = \mathcal{O}(1) \Rightarrow \|\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})\|_2 = \mathcal{O}(1)$. Since

$$\mathcal{L}_\lambda \left(\boldsymbol{\theta} \left(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right) \right) - \mathcal{L}_\lambda \left(\boldsymbol{\theta} \left(T_c^+(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}} \right) \right) = \int_{\frac{1}{\lambda} \log \alpha}^{T_c^+(\alpha)} \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}(t; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}))\|_2^2 dt$$

and $0 \leq \mathcal{L}_\lambda(\boldsymbol{\theta}(T_c^+(\alpha); \alpha \bar{\boldsymbol{\theta}}_{\text{init}})), \mathcal{L}_\lambda(\boldsymbol{\theta}(\frac{1}{\lambda} \log \alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}})) \leq \mathcal{O}(\lambda)$ by [Lemma D.1](#), we deduce that there exists $t_\alpha \in [\frac{1}{\lambda} \log \alpha, T_c^+(\alpha)]$ such that $\|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}(t_\alpha; \alpha \bar{\boldsymbol{\theta}}_{\text{init}}))\|_2 \leq \mathcal{O}(\sqrt{\lambda} \frac{c}{\lambda} \log \alpha) = \mathcal{O}(\lambda (\log \frac{1}{\lambda})^{-1/2})$.

For any sequence $\{\alpha_k\}_{k \geq 1}$ with $\alpha_k \rightarrow +\infty$, we choose $t_k = t_{\alpha_k}$. Let $\{\boldsymbol{\theta}(t_{i_k}; \alpha_{i_k} \bar{\boldsymbol{\theta}}_{\text{init}}) : k \geq 1\}$ be any convergent subsequence of $\{\boldsymbol{\theta}(t_k; \alpha_k \bar{\boldsymbol{\theta}}_{\text{init}}) : k \geq 1\}$, then we have that $\lim_{k \rightarrow +\infty} \|\nabla \mathcal{L}_\lambda(\boldsymbol{\theta}(t_{i_k}; \alpha_{i_k} \bar{\boldsymbol{\theta}}_{\text{init}}))\|_2 = 0$, i.e.,

$$\lim_{k \rightarrow +\infty} \left\| 2 \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}(t_{i_k}; \alpha_{i_k} \bar{\boldsymbol{\theta}}_{\text{init}})) + \lambda \boldsymbol{\theta}(t_{i_k}; \alpha_{i_k} \bar{\boldsymbol{\theta}}_{\text{init}}) \right\|_2 = 0. \quad (33)$$

By the choice of $\{i_k\}_{k \geq 1}$ we know that $\{\boldsymbol{\theta}(t_{i_k}; \alpha_{i_k} \bar{\boldsymbol{\theta}}_{\text{init}})\}_{k \geq 1}$ converges to a point $\boldsymbol{\theta}^* \in \mathbb{R}^d$, so the above implies that

$$2 \sum_{i=1}^n s_i \nabla f_i(\boldsymbol{\theta}^*) + \lambda \boldsymbol{\theta}^* = 0,$$

thus $\boldsymbol{\theta}^*$ is a KKT point of [\(R2\)](#), as desired. \square

E GENERALIZATION ANALYSES FOR LINEAR CLASSIFICATION

In this section, we provide supplementary generalization analyses for the linear classification experiments in Section 3.2.3. We follow the notations in Section 3.1. Let the input distribution \mathcal{D}_X , and the ground-truth label $y^*(\mathbf{x})$ be $\text{sign}(\mathbf{x}^\top \mathbf{w}^*)$ for an unknown $\mathbf{w}^* \in \mathbb{R}^d$. For a training set $S = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, we define $y_i := y^*(\mathbf{x}_i)$ for all $i \in [n]$.

E.1 PRELIMINARIES

We base our generalization analyses on Rademacher complexity.

Definition E.1 Given a family of functions \mathcal{F} mapping from \mathbb{R}^d to \mathbb{R} , the empirical Rademacher complexity of \mathcal{F} for a set of samples $S = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ is defined as

$$\hat{\mathfrak{R}}_S(\mathcal{F}) = \mathbb{E}_{\sigma_1, \dots, \sigma_n \sim \text{unif}\{\pm 1\}} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{x}_i) \right]. \quad (34)$$

Rademacher complexity is useful in deriving margin-based generalization bounds for classification. The following bound is standard in the literature Koltchinskii and Panchenko (2002); Mohri et al. (2018).

Theorem E.2 (Theorem 5.8, Mohri et al. (2018)) Let \mathcal{F} be a family of functions mapping from $\mathbb{R}^d \rightarrow \mathbb{R}$. $q > 0$. For any $\delta > 0$, with probability at least $1 - \delta$ over the random draw of $S = (\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mathcal{D}_X^n$, the following bound for the test error holds for all $f \in \mathcal{F}$:

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_X} [\mathbb{1}_{[y^*(\mathbf{x})f(\mathbf{x}) \leq 0]}] \leq \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{[y_i f(\mathbf{x}_i) \leq q]} + \frac{2}{q} \hat{\mathfrak{R}}_S(\mathcal{F}) + 3\sqrt{\frac{\log(2/\delta)}{2n}}. \quad (35)$$

Let $\mathcal{F}_p = \{\mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle : \|\mathbf{w}\|_p \leq 1\}$ be a family of linear functions on \mathbb{R}^d with bounded weight in L^p -norm. The following theorem from Awasthi et al. (2020) bounds the empirical Rademacher complexity of \mathcal{F}_p .

Theorem E.3 (Direct Corollary of Corollary 3 in Awasthi et al. (2020)) For a set of samples $S = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, the empirical Rademacher complexity of \mathcal{F}_1 and \mathcal{F}_2 is bounded as

$$\hat{\mathfrak{R}}_S(\mathcal{F}_1) \leq \frac{1}{n} \sqrt{2 \log(2d)} \cdot \left(\sum_{i=1}^n \|\mathbf{x}_i\|_\infty^2 \right)^{1/2}, \quad (36)$$

$$\hat{\mathfrak{R}}_S(\mathcal{F}_2) \leq \frac{1}{n} \left(\sum_{i=1}^n \|\mathbf{x}_i\|_2^2 \right)^{1/2}. \quad (37)$$

E.2 GENERALIZATION BOUNDS FOR SPARSE LINEAR CLASSIFICATION

Let \mathcal{D}_X be the uniform distribution on $\text{unif}\{\pm 1\}^d$. Let $k = \mathcal{O}(1)$ be a positive odd number. We draw \mathbf{w}^* as follows: randomly draw the first k coordinates from $\text{unif}\{\pm 1\}^k$, and set all the other coordinates to 0.

Theorem E.4 With probability at least $1 - \delta$ over the random draw of the training set $S = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, for any linear classifier $f : \mathbb{R}^d \rightarrow \mathbb{R}, \mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle$ that maximizes the L^1 -margin $\gamma_1(\mathbf{w}) := \min_{i \in [n]} \frac{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle}{\|\mathbf{w}\|_1}$ on S , the following bound for the test error holds:

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_X} [\mathbb{1}_{[y^*(\mathbf{x})f(\mathbf{x}) \leq 0]}] \leq 4k \cdot \sqrt{\frac{2 \log(2d)}{n}} + 3\sqrt{\frac{\log(2/\delta)}{2n}}. \quad (38)$$

Proof: It suffices to consider the case where $\|\mathbf{w}\|_1 = 1$ because rescaling \mathbf{w} does not change the test error. Since $\mathbf{w} \in \arg \max \gamma_1(\mathbf{w})$, $\gamma_1(\mathbf{w}) \geq \gamma_1(\mathbf{w}^*) = \frac{1}{k}$. By definition, this implies

$y_i f(\mathbf{x}_i) \geq \frac{1}{k}$. Letting $q := \frac{1}{2k}$ and applying [Theorem E.2](#), we have

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_X} [\mathbb{1}_{[y^*(\mathbf{x})f(\mathbf{x}) \leq 0]}] \leq 4k \cdot \hat{\mathfrak{R}}_S(\mathcal{F}_1) + 3\sqrt{\frac{\log(2/\delta)}{2n}}. \quad (39)$$

By [Theorem E.3](#), $\hat{\mathfrak{R}}_S(\mathcal{F}_1) \leq \frac{1}{n} \sqrt{2 \log(2d)} \cdot \sqrt{n} = \sqrt{\frac{2 \log(2d)}{n}}$. Combining this with (41) completes the proof. \square

Theorem E.5 *With probability at least $1 - \delta$ over the random draw of the training set $S = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, for any linear classifier $f : \mathbb{R}^d \rightarrow \mathbb{R}, \mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle$ that maximizes the L^2 -margin $\gamma_2(\mathbf{w}) := \min_{i \in [n]} \frac{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle}{\|\mathbf{w}\|_2}$ on S , the following bound for the test error holds:*

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_X} [\mathbb{1}_{[y^*(\mathbf{x})f(\mathbf{x}) \leq 0]}] \leq 4\sqrt{\frac{kd}{n}} + 3\sqrt{\frac{\log(2/\delta)}{2n}}. \quad (40)$$

Proof: It suffices to consider the case where $\|\mathbf{w}\|_2 = 1$ because rescaling \mathbf{w} does not change the test error. Since $\mathbf{w} \in \arg \max \gamma_2(\mathbf{w})$, $\gamma_2(\mathbf{w}) \geq \gamma_2(\mathbf{w}^*) = \frac{1}{\sqrt{k}}$. By definition, this implies $y_i f(\mathbf{x}_i) \geq \frac{1}{\sqrt{k}}$. Letting $q := \frac{1}{2\sqrt{k}}$ and applying [Theorem E.2](#), we have

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_X} [\mathbb{1}_{[y^*(\mathbf{x})f(\mathbf{x}) \leq 0]}] \leq 4\sqrt{k} \cdot \hat{\mathfrak{R}}_S(\mathcal{F}_2) + 3\sqrt{\frac{\log(2/\delta)}{2n}}. \quad (41)$$

By [Theorem E.3](#), $\hat{\mathfrak{R}}_S(\mathcal{F}_2) \leq \frac{1}{n} \cdot \sqrt{nd} = \sqrt{\frac{d}{n}}$. Combining this with (41) completes the proof. \square

F PROOFS FOR MATRIX COMPLETION

In this section, we give the detailed proof of results in [Section 3.3.3](#).

Let $\Omega = \{\mathbf{x}_k = (i_k, j_k) : 1 \leq k \leq n\}$ and we have n observations $\mathbf{P}_{\mathbf{x}_k} = \mathbf{e}_{i_k} \mathbf{e}_{j_k}^\top, 1 \leq k \leq n$. For convenience, we use \mathbf{P}_k to denote $\mathbf{P}_{\mathbf{x}_k}$. Define $\boldsymbol{\theta}_t = (\mathbf{U}_t, \mathbf{V}_t)$, $\mathbf{W}_t = \mathbf{U}_t \mathbf{U}_t^\top - \mathbf{V}_t \mathbf{V}_t^\top$ and $f_i(\boldsymbol{\theta}) = \langle \mathbf{W}_t, \mathbf{P}_i \rangle$, and our goal is to minimize the function (5). Since \mathbf{W}_t is symmetric, it is more convenient to replace \mathbf{P}_i with $\frac{1}{2}(\mathbf{P}_i + \mathbf{P}_i^\top)$. From now on, we let $\mathbf{P}_i = \frac{1}{2}(\mathbf{e}_{i_i} \mathbf{e}_{j_i}^\top + \mathbf{e}_{j_i} \mathbf{e}_{i_i}^\top)$.

The key step is to show that the loss \mathcal{L} defined in (5) satisfies the following two properties:

Property F.1 *All local minima of \mathcal{L} are global.*

Property F.2 *At any saddle point $(\mathbf{U}_s, \mathbf{V}_s)$ of \mathcal{L} , there is a direction $(\boldsymbol{\mathcal{E}}_U, \boldsymbol{\mathcal{E}}_V)$ such that*

$$\text{vec}(\boldsymbol{\mathcal{E}}_U, \boldsymbol{\mathcal{E}}_V)^\top \nabla^2 \mathcal{L}(\mathbf{U}_s, \mathbf{V}_s) \text{vec}(\boldsymbol{\mathcal{E}}_U, \boldsymbol{\mathcal{E}}_V) < 0.$$

Given these two properties, we can deduce that GF with random initialization can converge to global minimizers almost surely.

Theorem F.3 *The function \mathcal{L} satisfies [Properties F.1](#) and [F.2](#).*

Proof: Recall that

$$\mathcal{L}(\mathbf{U}, \mathbf{V}) = \frac{1}{n} \sum_{i=1}^n (\langle \mathbf{P}_i, \mathbf{U}\mathbf{U}^\top - \mathbf{V}\mathbf{V}^\top \rangle - y_i^*)^2 + \frac{\lambda}{2} \text{tr}(\mathbf{U}\mathbf{U}^\top + \mathbf{V}\mathbf{V}^\top).$$

Define

$$\hat{\mathcal{L}}(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n (\langle \mathbf{P}_i, \mathbf{W} \rangle - y_i^*)^2 + \frac{\lambda}{2} \|\mathbf{W}\|_*, \quad (42)$$

then we have $\mathcal{L}(\mathbf{U}, \mathbf{V}) \geq \hat{\mathcal{L}}(\mathbf{U}\mathbf{U}^\top - \mathbf{V}\mathbf{V}^\top)$. Since $\hat{\mathcal{L}}$ is convex, \mathbf{W} is a global minimizer of $\hat{\mathcal{L}}$ if and only if

$$0 \in \partial \nabla \hat{\mathcal{L}}(\mathbf{W}) = \frac{2}{n} \sum_{i=1}^n (\langle \mathbf{P}_i, \mathbf{W} \rangle - y_i^*) \mathbf{P}_i + \frac{\lambda}{2} \partial \|\mathbf{W}\|_*. \quad (43)$$

Moreover, for any $\mathbf{W} \in \mathbb{R}^{d \times d}$ it holds that

$$\min_{\mathbf{U}\mathbf{U}^\top - \mathbf{V}\mathbf{V}^\top = \mathbf{W}} \mathcal{L}(\mathbf{U}, \mathbf{V}) = \hat{\mathcal{L}}(\mathbf{W}). \quad (44)$$

now consider some $(\mathbf{U}_*, \mathbf{V}_*)$ such that

$$\nabla \mathcal{L}(\mathbf{U}_*, \mathbf{V}_*) = 0 \quad \text{and} \quad \nabla^2 \mathcal{L}(\mathbf{U}_*, \mathbf{V}_*) \succeq \mathbf{0}.$$

To prove the theorem's statement, we only need to show that $(\mathbf{U}_*, \mathbf{V}_*)$ is a global minimizer of \mathcal{L} . The subsequent proof is organized into two parts:

Claim 1. $\mathbf{W}^* = \mathbf{U}_* \mathbf{U}_*^\top - \mathbf{V}_* \mathbf{V}_*^\top$ is a global minimizer of $\hat{\mathcal{L}}$.

Proof of Claim 1: We check that the condition (43) holds at \mathbf{W}^* . First the first order condition imply that

$$\nabla_{\mathbf{U}} \mathcal{L}(\mathbf{U}_*, \mathbf{V}_*) = \left(\frac{4}{n} \sum_{i=1}^n (\langle \mathbf{P}_i, \mathbf{W}^* \rangle - y_i^*) \mathbf{P}_i + \lambda \mathbf{I} \right) \mathbf{U}_* = 0. \quad (45)$$

Similarly,

$$\left(-\frac{4}{n} \sum_{i=1}^n (\langle \mathbf{P}_i, \mathbf{W}^* \rangle - y_i^*) \mathbf{P}_i + \lambda \mathbf{I} \right) \mathbf{V}_* = 0. \quad (46)$$

Second, for any $\mathcal{E} \in \mathbb{R}^{d \times d}$ we have

$$\begin{aligned}
& (\nabla_U^2 \mathcal{L}(U_*, V_*)) \mathcal{E} = \lim_{t \rightarrow 0} \frac{1}{t} (\nabla_U \mathcal{L}(U_* + t\mathcal{E}, V_*) - \nabla_U \mathcal{L}(U_*, V_*)) \\
& = \lim_{t \rightarrow 0} \frac{1}{t} \frac{4}{n} \sum_{i=1}^n \left[\left(\langle (U_* + t\mathcal{E})(U_* + t\mathcal{E})^\top - V_* V_*^\top, P_i \rangle - y_i^* \right) P_i (U_* + t\mathcal{E}) \right. \\
& \quad \left. - (\langle P_i, W^* \rangle - y_i^*) P_i U_* \right] + \lambda \mathcal{E} \\
& = \frac{4}{n} \sum_{i=1}^n \left[(\langle P_i, W^* \rangle - y_i^*) P_i \mathcal{E} + 2 \langle U_* \mathcal{E}^\top, P_i \rangle P_i U_* \right] + \lambda \mathcal{E}.
\end{aligned} \tag{47}$$

where we use $\langle U_* \mathcal{E}^\top + \mathcal{E} U_*^\top, P_i \rangle = 2 \langle U_* \mathcal{E}^\top, P_i \rangle$ since P_i is symmetric. Similarly,

$$(\nabla_V^2 \mathcal{L}(U_*, V_*)) \mathcal{E} = \frac{4}{n} \sum_{i=1}^n \left[-(\langle P_i, W^* \rangle - y_i^*) P_i \mathcal{E} + 2 \langle V_* \mathcal{E}^\top, P_i \rangle P_i V_* \right] + \lambda \mathcal{E} \tag{48}$$

and

$$\nabla_V \nabla_U \mathcal{L}(U_*, V_*) \mathcal{E} = -\frac{8}{n} \sum_{i=1}^n \langle V_* \mathcal{E}^\top, P_i \rangle P_i U_* \tag{49}$$

Let $M = \frac{4}{n} \sum_{i=1}^n (\langle P_i, W^* \rangle - y_i^*) P_i$. Since $\nabla^2 \mathcal{L}(U_*, V_*) \succeq \mathbf{0}$, for any $\mathcal{E}_U, \mathcal{E}_V \in \mathbb{R}^{d \times d}$ we have

$$\begin{aligned}
& 0 \leq \langle \mathcal{E}_U, (\nabla_U^2 \mathcal{L}(U_*, V_*)) \mathcal{E}_U \rangle + \langle \mathcal{E}_V, (\nabla_V^2 \mathcal{L}(U_*, V_*)) \mathcal{E}_V \rangle + 2 \langle \mathcal{E}_U, \nabla_V \nabla_U \mathcal{L}(U_*, V_*) \mathcal{E}_V \rangle \\
& = \langle \mathcal{E}_U, (M + \lambda I) \mathcal{E}_U \rangle + \langle \mathcal{E}_V, (-M + \lambda I) \mathcal{E}_V \rangle - \frac{16}{n} \sum_{i=1}^n \langle U_* \mathcal{E}_U^\top, P_i \rangle \langle V_* \mathcal{E}_V^\top, P_i \rangle \\
& \quad + \frac{8}{n} \sum_{i=1}^n \left(\langle V_* \mathcal{E}_V^\top, P_i \rangle^2 + \langle U_* \mathcal{E}_U^\top, P_i \rangle^2 \right).
\end{aligned} \tag{50}$$

We now consider two cases:

- $\max \{\text{rank} U_*, \text{rank} V_*\} = d$. We assume WLOG that $\text{rank} U_* = d$, then (45) immediately implies that $M + \lambda I = 0$. Then $-M + \lambda I = 2\lambda I$ and by (46) we have $V_* = 0$. In this case, W^* is positive semi-definite, and (43) holds since $I \in \partial \|\mathbf{W}^*\|_*$ by Lemma F.4.
- $\max \{\text{rank} U_*, \text{rank} V_*\} < d$. In this case, there exists non-zero vectors $\mathbf{a}, \mathbf{c} \in \mathbb{R}^d$ such that $U_* \mathbf{a} = V_* \mathbf{c} = 0$. For arbitrary vectors $\mathbf{b}, \mathbf{d} \in \mathbb{R}^d$, we set $\mathcal{E}_U = \mathbf{b} \mathbf{a}^\top$ and $\mathcal{E}_V = \mathbf{d} \mathbf{c}^\top$ in (50), so that

$$\|\mathbf{a}\|^2 \mathbf{b}^\top (M + \lambda I) \mathbf{b} + \|\mathbf{c}\|^2 \mathbf{d}^\top (\lambda I - M) \mathbf{d} \geq 0.$$

Since \mathbf{b}, \mathbf{d} are arbitrarily chosen, the above implies that

$$M + \lambda I \succeq 0, \quad \text{and} \quad \lambda I - M \succeq 0. \tag{51}$$

Since M is symmetric, there exists an orthogonal basis $\{\mathbf{f}_i : i \in [d]\}$ of \mathbb{R}^d and corresponding eigenvalues $m_i, i \in [d]$ such that $M \mathbf{f}_i = m_i \mathbf{f}_i$. We partition the set $[d]$ into three subsets:

$$\mathcal{S}_- = \{i \in [d] : m_i = -\lambda\}, \quad \mathcal{S}_+ = \{i \in [d] : m_i = \lambda\}, \quad \mathcal{S}_0 = [d] - \mathcal{S}_- - \mathcal{S}_+.$$

Since $(M + \lambda I) U_* = 0$, we have $U_*^\top \mathbf{f}_i = 0$ for all $i \in \mathcal{S}_+ \cup \mathcal{S}_0$. As a result, we can write

$$U_* U_*^\top = \sum_{i,j \in \mathcal{S}_-} a_{ij} \mathbf{f}_i \mathbf{f}_j^\top, \quad a_{ij} \in \mathbb{R}.$$

Similarly, we can write

$$V_* V_*^\top = \sum_{i,j \in \mathcal{S}_+} a_{ij} \mathbf{f}_i \mathbf{f}_j^\top, \quad a_{ij} \in \mathbb{R}.$$

Assume WLOG that $\mathcal{S}_- = [t]$ and $\mathcal{S}_+ = \{t+1, t+2, \dots, s\}$, then Hence $\mathbf{W}^* = \sum_{i,j=1}^t a_{ij} \mathbf{f}_i \mathbf{f}_j^\top - \sum_{i,j=t+1}^s a_{ij} \mathbf{f}_i \mathbf{f}_j^\top$. note that the matrices $(a_{ij})_{i,j \in \mathcal{S}_-}$ and $(a_{ij})_{i,j \in \mathcal{S}_+}$ are both positive semi-definite, so there exists $b_i \geq 0, i \in [d]$ and an orthogonal basis $\{\mathbf{g}_i : i \in [d]\}$ such that

$$\sum_{i,j=1}^t a_{ij} \mathbf{f}_i \mathbf{f}_j^\top = \sum_{i=1}^t b_i \mathbf{g}_i \mathbf{g}_i^\top \quad \text{and} \quad \sum_{i,j=t+1}^s a_{ij} \mathbf{f}_i \mathbf{f}_j^\top = \sum_{i=t+1}^s b_i \mathbf{g}_i \mathbf{g}_i^\top \quad (52)$$

and $b_i = 0$ and $\mathbf{g}_i = \mathbf{f}_i$ for $i > s$. now let

$$\mathbf{G}_1 = (\mathbf{g}_i : i \in [d]), \quad \mathbf{G}_2 = \left((-1)^{\mathbb{I}\{i>t\}} \mathbf{g}_i : i \in [d] \right), \quad \text{and} \quad \mathbf{B} = \text{diag}\{b_i : i \in [d]\},$$

then both \mathbf{G}_1 and \mathbf{G}_2 are orthonormal matrices and we have $\mathbf{W}^* = \mathbf{G}_1 \mathbf{B} \mathbf{G}_2^\top$. This gives a singular value decomposition of \mathbf{W}^* . By [Lemma F.4](#) we can write down the explicit form of $\partial \|\mathbf{W}^*\|_*$ as follows:

$$\partial \|\mathbf{W}^*\|_* = \left\{ \sum_{i \in \mathcal{T}} (-1)^{\mathbb{I}\{i>t\}} \mathbf{g}_i \mathbf{g}_i^\top + \mathbf{E} : \|\mathbf{E}\| \leq 1 \text{ and } \mathbf{E} \mathbf{g}_i = 0, \forall i \in \mathcal{T} \right\}.$$

Here $\mathcal{T} = \{i \in [d] : b_i \neq 0\} \subset [s]$. We choose

$$\begin{aligned} \mathbf{E} &= - \left(\lambda^{-1} \mathbf{M} + \sum_{i \in \mathcal{T}} (-1)^{\mathbb{I}\{i>t\}} \mathbf{g}_i \mathbf{g}_i^\top \right) \\ &= -\lambda^{-1} \left(-\lambda \sum_{i=1}^t \mathbf{f}_i \mathbf{f}_i^\top + \lambda \sum_{i=t+1}^s \mathbf{f}_i \mathbf{f}_i^\top + \sum_{i>s} m_i \mathbf{f}_i \mathbf{f}_i^\top \right) - \sum_{i \in \mathcal{T}} (-1)^{\mathbb{I}\{i>t\}} \mathbf{g}_i \mathbf{g}_i^\top \\ &= -\lambda^{-1} \left(-\lambda \sum_{i=1}^t \mathbf{g}_i \mathbf{g}_i^\top + \lambda \sum_{i=t+1}^s \mathbf{g}_i \mathbf{g}_i^\top + \sum_{i>s} m_i \mathbf{g}_i \mathbf{g}_i^\top \right) - \sum_{i \in \mathcal{T}} (-1)^{\mathbb{I}\{i>t\}} \mathbf{g}_i \mathbf{g}_i^\top \\ &= \sum_{i \in [s] - \mathcal{T}} (-1)^{\mathbb{I}\{i>t\}} \mathbf{g}_i \mathbf{g}_i^\top - \lambda^{-1} \sum_{i>s} m_i \mathbf{g}_i \mathbf{g}_i^\top \end{aligned}$$

where the third equation holds because of the definition (52) and $\mathbf{g}_i = \mathbf{f}_i$ when $i > s$. Since by (51) we have $m_i \in [-\lambda, \lambda]$ for all $i \in [d]$, the above expression of \mathbf{E} immediately implies that $\|\mathbf{E}\| \leq 1$. Moreover, we obviously have $\mathbf{E} \mathbf{g}_i = 0$ when $i \in \mathcal{T}$ (since $\{\mathbf{g}_i\}$ is orthogonal). Hence, we have $0 \in \mathbf{M} + \lambda \partial \|\mathbf{W}^*\|_*$.

To summarize, we have shown that (43) always holds at \mathbf{W}^* . Therefore, \mathbf{W}^* is a global minimizer of $\hat{\mathcal{L}}$ which concludes the proof of [Claim 1](#).

Claim 2. $(\mathbf{U}_*, \mathbf{V}_*)$ solves the problem

$$\text{minimize} \quad \|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2 \quad \text{s.t.} \quad \mathbf{U} \mathbf{U}^\top - \mathbf{V} \mathbf{V}^\top = \mathbf{W}^*. \quad (53)$$

Proof of Claim 2. We will use the notations in the proof of Claim 1 for convenience. note that

$$\|\mathbf{U}_*\|_F^2 = \text{tr} \mathbf{U}_* \mathbf{U}_*^\top = \sum_{i=1}^t b_i$$

and similarly

$$\|\mathbf{V}_*\|_F^2 = \sum_{i=t+1}^s b_i.$$

Thus $\|\mathbf{U}_*\|_F^2 + \|\mathbf{V}_*\|_F^2 = \sum_{i=1}^s b_i = \|\mathbf{W}^*\|_*$. On the other hand, for any (\mathbf{U}, \mathbf{V}) satisfying $\mathbf{U} \mathbf{U}^\top - \mathbf{V} \mathbf{V}^\top = \mathbf{W}^*$, we have

$$\|\mathbf{W}^*\|_* = \|\mathbf{U} \mathbf{U}^\top - \mathbf{V} \mathbf{V}^\top\|_* \leq \|\mathbf{U} \mathbf{U}^\top\|_* + \|\mathbf{V} \mathbf{V}^\top\|_* = \|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2.$$

Hence $(\mathbf{U}_*, \mathbf{V}_*)$ is a minimizer of (53), as desired.

We can now prove that $(\mathbf{U}_*, \mathbf{V}_*)$ is a global minimizer of \mathcal{L} , which completes the proof of **Theorem F.3**. Indeed, for any (\mathbf{U}, \mathbf{V}) we have

$$\mathcal{L}(\mathbf{U}, \mathbf{V}) \geq \hat{\mathcal{L}}(\mathbf{U}\mathbf{U}^\top - \mathbf{V}\mathbf{V}^\top) \geq \hat{\mathcal{L}}(\mathbf{W}^*) = \mathcal{L}(\mathbf{U}_*, \mathbf{V}_*).$$

The last step follows from **Claim 2** and (44). \square

Lemma F.4 (Watson, 1992, Example 2) *Let $\mathbf{A} = \mathbf{L}\mathbf{\Sigma}\mathbf{R}^\top$ be its singular value decomposition, then we have*

$$\partial\|\mathbf{A}\|_* = \{\mathbf{L}\mathbf{R}^\top + \mathbf{E} : \|\mathbf{E}\| \leq 1, \mathbf{L}^\top \mathbf{E} = 0 \text{ and } \mathbf{E}\mathbf{R} = 0\}.$$

It is well-known that under certain regularity conditions on the ground-truth matrix \mathbf{X}^* , the global minimizer of the convex problem (42) is close to the ground-truth \mathbf{X}^* . Here we present a version of this result adapted from (Candes and Plan, 2010, Theorem 7).

Theorem F.5 *Suppose that $\text{rank}(\mathbf{X}^*) = r = \mathcal{O}(1)$, $\mathbf{X}^* = \mathbf{V}_{\mathbf{X}^*}\mathbf{\Sigma}_{\mathbf{X}^*}\mathbf{V}_{\mathbf{X}^*}^\top$ is (a version of) its SVD, and each row of $\mathbf{V}_{\mathbf{X}^*}$ has ℓ_∞ -norm bounded by $\sqrt{\frac{\mu}{d}}$, then if the number of observed entries satisfies $n \gtrsim \mu^4 d \log^2 d$, then we have that $\|\mathbf{W}^* - \mathbf{X}^*\|_F \lesssim \sqrt{\lambda\|\mathbf{X}^*\|_*}\mu^2 \log d$ with probability $\geq 1 - d^{-3}$.*