

A Unifying View on Implicit Bias in Training Linear Neural Networks

Chulhee Yun*

MIT

CHULHEEY@MIT.EDU

Shankar Krishnan

Hossein Mobahi

Google Research

SKRISHNAN@GOOGLE.COM

HMOBAHI@GOOGLE.COM

Abstract

We study the implicit bias of gradient flow (i.e., gradient descent with infinitesimal step size) on linear neural network training. We propose a tensor formulation of neural networks that includes fully-connected, diagonal, and convolutional networks as special cases, and investigate the linear version of the formulation called linear tensor networks. With this formulation, we can identify the convergence direction of the network parameters as singular vectors of a tensor defined by the network. For L -layer linear tensor networks that are orthogonally decomposable, we show that gradient flow on separable classification finds a stationary point of the $\ell_{2/L}$ max-margin problem in a “transformed” input space defined by the network. For underdetermined regression, we prove that gradient flow finds a global minimum which minimizes a norm-like function that interpolates between weighted ℓ_1 and ℓ_2 norms in the transformed input space. Our theorems subsume existing results in the literature while removing standard convergence assumptions.

1. Introduction

Overparametrized neural networks have infinitely many solutions that achieve zero training error, and such global minima have different generalization performance. Moreover, training a neural network is a high-dimensional nonconvex problem, which is typically intractable to solve. However, the success of deep learning indicates that first-order methods such as gradient descent or stochastic gradient descent (GD/SGD) not only (a) succeed in finding global minima, but also (b) are biased towards solutions that generalize well, which largely has remained a mystery in the literature.

To explain part (a) of the phenomenon, there is a growing literature studying the convergence of GD/SGD on overparametrized neural networks (e.g., [1, 9, 10, 17, 28, 33]). There are also convergence results that focus on linear networks [3, 5, 8, 16, 31]. These results typically focus on the convergence of loss, hence do not address *which* of the many global minima is reached.

Another line of results tackles part (b), by studying the implicit bias or regularization of gradient-based methods on neural networks or related problems [2, 4, 6, 11–14, 18, 29, 30]. These results have shown interesting progress that even without explicit regularization terms in the training objective, algorithms such as GD applied on neural networks have an *implicit bias* towards certain solutions among the many global minima. However, in proving such results, many results rely on *convergence assumptions* such as global convergence of loss to zero and/or directional convergence of parameters and gradients, which ideally should be removed. We cover more details on this literature in Section B.

* Based on work performed during internship at Google Research

1.1. Summary of our contributions

We study the implicit bias of gradient flow (GD with infinitesimal step size) on linear neural networks. Following recent literature, we consider classification and regression problems that have multiple solutions with zero training error. **Our analyses apply to a general class of networks, and prove both convergence and bias, overcoming the limitations of the existing results.**

We propose a general *tensor formulation* of nonlinear neural networks which includes many network architectures considered in the literature. In this paper, we focus on the linear version of this formulation, called *linear tensor networks*. For linearly separable classification, we prove that the network parameters converge in direction to *singular vectors* of a tensor defined by the network. Also, we show that if the linear tensor network is orthogonally decomposable (Assumption 1), the gradient flow finds the $\ell_{2/\text{depth}}$ max-margin solution in the singular value space, leading the parameters to converge to the *top singular vectors* of the tensor when depth = 2. For underdetermined linear regression, we study the limit points of gradient flow on orthogonally decomposable networks (Assumption 1), and provide a full characterization of the limit points. Due to space limits, additional theorems on 2-layer linear tensor networks in an extremely overparametrized setting, as well as results for specific network architectures (fully-connected, diagonal, and convolutional), are deferred to Appendix D and E. We also present our simple experiments in Appendix F.

2. Problem settings and tensor formulation of neural networks

Given a positive integer a , let $[a] := \{1, \dots, a\}$. We use \mathbf{I}_d to denote the $d \times d$ identity matrix. For two vectors \mathbf{a} and \mathbf{b} , let $\mathbf{a} \otimes \mathbf{b}$ be their tensor product, $\mathbf{a} \odot \mathbf{b}$ be their element-wise product, and $\mathbf{a}^{\odot k}$ be the element-wise k -th power of \mathbf{a} . Given an order- L tensor $\mathbf{A} \in \mathbb{R}^{k_1 \times \dots \times k_L}$, we use $[\mathbf{A}]_{j_1, \dots, j_L}$ to denote (j_1, j_2, \dots, j_L) -th element of \mathbf{A} , where $j_l \in [k_l]$ for all $l \in [L]$. In element indexing, we use \cdot to denote all indices in the corresponding dimension. We use \mathbf{e}_j^k to denote the j -th standard basis vector of the vector space \mathbb{R}^k . Given a tensor $\mathbf{A} \in \mathbb{R}^{k_1 \times \dots \times k_L}$ and linear maps $\mathbf{B}_l \in \mathbb{R}^{p_l \times k_l}$ for $l \in [L]$, we define the multilinear multiplication \circ between them as

$$\begin{aligned} \mathbf{A} \circ (\mathbf{B}_1^T, \mathbf{B}_2^T, \dots, \mathbf{B}_L^T) &= \sum_{j_1, \dots, j_L} [\mathbf{A}]_{j_1, \dots, j_L} (\mathbf{e}_{j_1}^{k_1} \otimes \dots \otimes \mathbf{e}_{j_L}^{k_L}) \circ (\mathbf{B}_1^T, \dots, \mathbf{B}_L^T) \\ &:= \sum_{j_1, \dots, j_L} [\mathbf{A}]_{j_1, \dots, j_L} (\mathbf{B}_1 \mathbf{e}_{j_1}^{k_1} \otimes \dots \otimes \mathbf{B}_L \mathbf{e}_{j_L}^{k_L}) \in \mathbb{R}^{p_1 \times \dots \times p_L}. \end{aligned}$$

2.1. Problem settings

We are given a dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$. We let $\mathbf{X} \in \mathbb{R}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}^n$ be the data matrix and label vector, respectively. We study binary classification and linear regression in this paper, in order to focus on the settings where there exist *many* global solutions. For binary classification, we assume $y_i \in \{\pm 1\}$ and that the data is separable: there exists a unit vector \mathbf{z} such that $y_i \mathbf{x}_i^T \mathbf{z} \geq \gamma > 0$ for all $i \in [n]$. For regression, we consider the underdetermined case ($n \leq d$) where there are many parameters $\mathbf{z} \in \mathbb{R}^d$ such that $\mathbf{X}\mathbf{z} = \mathbf{y}$. We assume \mathbf{X} has full row rank.

We use $f(\cdot; \Theta) : \mathbb{R}^d \rightarrow \mathbb{R}$ to denote a model parametrized by Θ . Given the network and the dataset, we consider minimizing the training loss $\mathcal{L}(\Theta) := \sum_{i=1}^n \ell(f(\mathbf{x}_i; \Theta), y_i)$ over Θ . Following previous results (e.g., [21, 23]), we use the exponential loss $\ell(\hat{y}, y) = \exp(-\hat{y}y)$ for classification problems. For regression, we use the squared error loss $\ell(\hat{y}, y) = \frac{1}{2}(\hat{y} - y)^2$. On the algorithm side, we minimize \mathcal{L} using gradient flow, which can be viewed as GD with infinitesimal step size. The gradient flow dynamics is defined as $\frac{d}{dt} \Theta = -\nabla_{\Theta} \mathcal{L}(\Theta)$.

2.2. Tensor formulation of neural networks

We now present a general tensor formulation of neural networks. Given an input $\mathbf{x} \in \mathbb{R}^d$, the network uses a linear map \mathbf{M} that maps \mathbf{x} to a order- L tensor $\mathbf{M}(\mathbf{x}) \in \mathbb{R}^{k_1 \times \dots \times k_L}$, where $L \geq 2$. Using parameters $\mathbf{v}_l \in \mathbb{R}^{k_l}$ and activation ϕ , the network computes its layers as the following:

$$\begin{aligned} \mathbf{H}_1(\mathbf{x}) &= \phi(\mathbf{M}(\mathbf{x}) \circ (\mathbf{v}_1, \mathbf{I}_{k_2}, \dots, \mathbf{I}_{k_L})) \in \mathbb{R}^{k_2 \times \dots \times k_L}, \\ \mathbf{H}_l(\mathbf{x}) &= \phi(\mathbf{H}_{l-1}(\mathbf{x}) \circ (\mathbf{v}_l, \mathbf{I}_{k_{l+1}}, \dots, \mathbf{I}_{k_L})) \in \mathbb{R}^{k_{l+1} \times \dots \times k_L}, \quad \text{for } l = 2, \dots, L-1, \\ f(\mathbf{x}; \Theta) &= \mathbf{H}_{L-1}(\mathbf{x}) \circ \mathbf{v}_L \in \mathbb{R}. \end{aligned} \quad (1)$$

We use Θ to denote the collection of all parameters $(\mathbf{v}_1, \dots, \mathbf{v}_L)$. We call $\mathbf{M}(\mathbf{x})$ the *data tensor*. This new formulation is general enough to capture many architectures considered in the literature, including **fully-connected networks, diagonal networks, and circular convolutional networks**. In Appendix C, we formally define these architectures and show that they are special cases of (1).

For the remaining of the paper, we study the case whose activation is *linear*, i.e., $\phi(t) = t$:

$$f(\mathbf{x}; \Theta) = \mathbf{M}(\mathbf{x}) \circ (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_L). \quad (2)$$

We will refer to (2) as *linear tensor networks*. Since the data tensor $\mathbf{M}(\mathbf{x})$ is a linear function of \mathbf{x} , the linear tensor network is also a linear function of \mathbf{x} . Thus, the output of the network can also be written as $f(\mathbf{x}; \Theta) = \mathbf{x}^T \beta(\Theta)$, where $\beta(\Theta) \in \mathbb{R}^d$ denotes the linear coefficients computed as a function of the network parameters Θ . Note that the expressive power of any linear tensor network is at best a linear model $\mathbf{x} \mapsto \mathbf{x}^T \mathbf{z}$; however, even though the models have the same expressive power, their architectural differences lead to different implicit biases in training, which we investigate in this paper. Studying separable classification and underdetermined regression is useful for highlighting such biases because there are *infinitely many* coefficients that perfectly classify or fit the dataset.

For our linear tensor network, the evolution of the parameters \mathbf{v}_l via gradient flow reads

$$\begin{aligned} \dot{\mathbf{v}}_l &= -\nabla_{\mathbf{v}_l} \mathcal{L}(\Theta) = -\sum_{i=1}^n \ell'(f(\mathbf{x}_i; \Theta), y_i) \mathbf{M}(\mathbf{x}_i) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L) \\ &= \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L), \quad \forall l \in [L], \end{aligned}$$

where we initialize $\mathbf{v}_l(0) = \alpha \bar{\mathbf{v}}_l$, for $l \in [L]$. We refer to α and $\bar{\mathbf{v}}_l$ as the *initial scale* and *initial direction*, respectively. The vector $\mathbf{r} \in \mathbb{R}^n$ is the residual vector, defined as

$$[\mathbf{r}]_i = \ell'(f(\mathbf{x}_i; \Theta), y_i) = \begin{cases} -y_i \exp(-y_i f(\mathbf{x}_i; \Theta)) & \text{for classification,} \\ f(\mathbf{x}_i; \Theta) - y_i & \text{for regression.} \end{cases} \quad (3)$$

3. Implicit bias of gradient flow on separable classification

In this section, we present our results on the implicit bias of gradient flow in binary classification with linearly separable data. Consider the singular value decomposition (SVD) of a matrix $\mathbf{A} = \sum_{j=1}^m s_j (\mathbf{u}_j \otimes \mathbf{v}_j)$, where m is the rank of \mathbf{A} . Note that the tuples $(\mathbf{u}_j, \mathbf{v}_j, s_j)$ are solutions to the system of equations $s\mathbf{u} = \mathbf{A}\mathbf{v}$ and $s\mathbf{v} = \mathbf{A}^T \mathbf{u}$. Lim [22] generalizes this definition of singular vectors and singular values to higher-order tensors: given an order- L tensor $\mathbf{A} \in \mathbb{R}^{k_1 \times \dots \times k_L}$, we define the singular vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_L$ and singular value s to be the solution of the following:

$$s\mathbf{u}_l = \mathbf{A} \circ (\mathbf{u}_1, \dots, \mathbf{u}_{l-1}, \mathbf{I}_{k_l}, \mathbf{u}_{l+1}, \dots, \mathbf{u}_L), \quad \text{for } l \in [L]. \quad (4)$$

Using the definition of the singular vectors of tensors, we can characterize the limit direction of parameters after reaching 100% training accuracy. In Appendix G, we prove the following:

Theorem 1 *If the gradient flow satisfies $\mathcal{L}(\Theta(t_0)) < 1$ for some $t_0 \geq 0$ and $\mathbf{X}^T \mathbf{r}(t)$ converges in direction, then $\mathbf{v}_1, \dots, \mathbf{v}_L$ converge to the singular vectors of $\mathbf{M}(-\lim_{t \rightarrow \infty} \mathbf{X}^T \mathbf{r}(t) / \|\mathbf{X}^T \mathbf{r}(t)\|)$.*

For this theorem, we make some convergence assumptions, because the network is fully general; this is the *only* result where we assume convergence. In fact, for the special case of linear fully-connected networks, the directional convergence assumption is *not* required, and their linear coefficients converge in direction to the ℓ_2 max-margin classifier. We state this corollary in Appendix E.2; this result also appears in Ji and Telgarsky [21], but we provide an alternative proof.

Admittedly, Theorem 1 is not a *full* characterization of the limit directions, because there are usually multiple solutions that satisfy (4). However, a better characterization of the directions is difficult, because singular vectors of high-order tensors are not well-understood as well as intractable to compute [15]. Given this intractability, it might be reasonable to make some assumptions on the “structure” of the data tensor $\mathbf{M}(\mathbf{x})$, so that they are easier to handle. The following assumption defines a class of *orthogonally decomposable* data tensors, which includes **linear diagonal networks** and **linear full-length convolutional networks** as special cases (see Appendix H.2 and H.3).

Assumption 1 *For the data tensor $\mathbf{M}(\mathbf{x}) \in \mathbb{R}^{k_1 \times \dots \times k_L}$ of a linear tensor network (2), there exist a full column rank matrix $\mathbf{S} \in \mathbb{C}^{m \times d}$ ($d \leq m \leq \min_l k_l$) and matrices $\mathbf{U}_1 \in \mathbb{C}^{k_1 \times m}, \dots, \mathbf{U}_L \in \mathbb{C}^{k_L \times m}$ such that $\mathbf{U}_l^H \mathbf{U}_l = \mathbf{I}_{k_l}$ for all $l \in [L]$, and the data tensor $\mathbf{M}(\mathbf{x})$ can be written as*

$$\mathbf{M}(\mathbf{x}) = \sum_{j=1}^m [\mathbf{S}\mathbf{x}]_j ([\mathbf{U}_1]_{\cdot,j} \otimes [\mathbf{U}_2]_{\cdot,j} \otimes \dots \otimes [\mathbf{U}_L]_{\cdot,j}). \quad (5)$$

In this assumption, we allow $\mathbf{U}_1, \dots, \mathbf{U}_L$ and \mathbf{S} to be complex matrices, although $\mathbf{M}(\mathbf{x})$ and parameters \mathbf{v}_l stay real, as defined earlier. For a complex matrix \mathbf{A} , we use \mathbf{A}^* to denote its entry-wise complex conjugate, \mathbf{A}^T to denote its transpose (without conjugating), and \mathbf{A}^H to denote its conjugate transpose. Using the assumption, the following theorem characterizes the limit directions.

Theorem 2 *Suppose a linear tensor network satisfies Assumption 1. If there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{v}}_1, \dots, \bar{\mathbf{v}}_L$ of the network parameters satisfy $|[\mathbf{U}_l^T \bar{\mathbf{v}}_l]_j|^2 - |[\mathbf{U}_L^T \bar{\mathbf{v}}_L]_j|^2 \geq \lambda$ for all $l \in [L-1]$ and $j \in [m]$, then $\beta(\Theta(t))$ converges in a direction that aligns with $\mathbf{S}^T \boldsymbol{\rho}^\infty$, where $\boldsymbol{\rho}^\infty \in \mathbb{C}^m$ denotes a stationary point of the following optimization problem*

$$\text{minimize}_{\boldsymbol{\rho} \in \mathbb{C}^m} \quad \|\boldsymbol{\rho}\|_{2/L} \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{S}^T \boldsymbol{\rho} \geq 1, \quad \forall i \in [n].$$

If \mathbf{S} is invertible, then $\beta(\Theta(t))$ converges in a direction that aligns with a stationary point \mathbf{z}^∞ of

$$\text{minimize}_{\mathbf{z} \in \mathbb{R}^d} \quad \|\mathbf{S}^{-T} \mathbf{z}\|_{2/L} \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{z} \geq 1, \quad \forall i \in [n].$$

Theorem 2 shows that the gradient flow finds sparse $\boldsymbol{\rho}^\infty$ that minimizes the $\ell_{2/L}$ norm in the “singular value space,” where the data points \mathbf{x}_i are transformed into the singular values $\mathbf{S}\mathbf{x}_i$. Also, the proof of Theorem 2 reveals that in case of $L = 2$, the parameters $\mathbf{v}_l(t)$ in fact converge to the *top* singular vectors of the data tensor $\mathbf{M}(-\mathbf{X}^T \mathbf{r})$; thus, compared to Theorem 1, we have a more complete characterization of “which” singular vectors to converge to. The proof is in Appendix H.

In our theorems, we pose assumptions on initial directions $\bar{\mathbf{v}}_l$ that are sufficient conditions for the loss $\mathcal{L}(\Theta(t))$ to converge to zero. Note that λ can be arbitrarily small, so the condition is satisfied **with probability 1** if we set $\bar{\mathbf{v}}_L = \mathbf{0}$ and sample $\bar{\mathbf{v}}_l$ ’s from i.i.d. Gaussian distribution.

Theorem 2 leads to corollaries (stated in Appendix E.3) on linear diagonal and full-length convolutional networks, showing that diagonal (or convolutional) networks converge to the stationary point of the max-margin problem with respect to the $\ell_{2/L}$ norm (or DFT-domain $\ell_{2/L}$ norm). Although this was already shown in Gunasekar et al. [14], note that Theorem 2 does not rely on convergence assumptions, which is different from the prior results.

4. Implicit bias of gradient flow on underdetermined regression

In Section 3, we observed that the limit directions of parameters do not depend on initialization. This is not the case in regression setting, because parameters do not diverge to infinity. As we show in this section, the limit points are closely tied to initialization. For the orthogonally decomposable networks (Assumption 1) with real \mathbf{S} and \mathbf{U}_l 's, we show how limit points of gradient flow change according to initialization. We consider an initialization scheme that, in the special case of diagonal networks, corresponds to initializing the first $L - 1$ layers to the same value and the last layer to zero. We use the following lemma on a relevant system of ODEs:

Lemma 3 *Consider the system of ODEs, where $p, q : \mathbb{R} \rightarrow \mathbb{R}$:*

$$\dot{p} = p^{L-2}q, \quad \dot{q} = p^{L-1}, \quad p(0) = 1, \quad q(0) = 0.$$

Then, the solutions $p_L(t)$ and $q_L(t)$ are continuous on their maximal interval of existence of the form $(-c, c) \subset \mathbb{R}$ for some $c \in (0, \infty]$. Define $h_L(t) = p_L(t)^{L-1}q_L(t)$; then, $h_L(t)$ is odd and strictly increasing, satisfying $\lim_{t \uparrow c} h_L(t) = \infty$ and $\lim_{t \downarrow -c} h_L(t) = -\infty$.

Using the function $h_L(t)$ from Lemma 3, we can obtain the following theorem that characterizes the limit points as the minimizer of a norm-like function $Q_{L,\alpha,\bar{\eta}}$ among the global minima.

Theorem 4 *Suppose a linear tensor network satisfies Assumption 1. Assume further that the matrices $\mathbf{U}_1, \dots, \mathbf{U}_L$ and \mathbf{S} from Assumption 1 are all real. For some $\lambda > 0$, choose any vector $\bar{\eta} \in \mathbb{R}^m$ satisfying $[\bar{\eta}]_j^2 \geq \lambda$ for all $j \in [m]$, and choose initial directions $\bar{\mathbf{v}}_l = \mathbf{U}_l \bar{\eta}$ for $l \in [L - 1]$ and $\bar{\mathbf{v}}_L = \mathbf{0}$. Then, the linear coefficients $\beta(\Theta(t))$ converge to $\mathbf{S}^T \boldsymbol{\rho}^\infty$, where $\boldsymbol{\rho}^\infty$ is the solution of*

$$\text{minimize}_{\boldsymbol{\rho} \in \mathbb{R}^m} \quad Q_{L,\alpha,\bar{\eta}}(\boldsymbol{\rho}) := \alpha^2 \sum_{j=1}^m [\bar{\eta}]_j^2 H_L \left(\frac{[\boldsymbol{\rho}]_j}{\alpha^L |[\bar{\eta}]_j|^L} \right) \quad \text{subject to} \quad \mathbf{X} \mathbf{S}^T \boldsymbol{\rho} = \mathbf{y},$$

where $Q_{L,\alpha,\bar{\eta}} : \mathbb{R}^m \rightarrow \mathbb{R}$ is a norm-like function defined using $H_L(t) := \int_0^t h_L^{-1}(\tau) d\tau$. If \mathbf{S} is invertible, then $\beta(\Theta(t))$ converges to the solution \mathbf{z}^∞ of

$$\text{minimize}_{\mathbf{z} \in \mathbb{R}^d} \quad Q_{L,\alpha,\bar{\eta}}(\mathbf{S}^{-T} \mathbf{z}) \quad \text{subject to} \quad \mathbf{X} \mathbf{z} = \mathbf{y}.$$

The proofs of Lemma 3 and Theorem 4 are deferred to Appendix I. It can be checked that $H_L(t)$ grows like the absolute value function if t is large, and grows like a quadratic function if t is close to zero. This means that

$$\lim_{\alpha \rightarrow 0} Q_{L,\alpha,\bar{\eta}}(\boldsymbol{\rho}) \propto \sum_{j=1}^m \frac{|[\boldsymbol{\rho}]_j|}{|[\bar{\eta}]_j|^{L-2}}, \quad \lim_{\alpha \rightarrow \infty} Q_{L,\alpha,\bar{\eta}}(\boldsymbol{\rho}) \propto \sum_{j=1}^m \frac{[\boldsymbol{\rho}]_j^2}{[\bar{\eta}]_j^{2L-2}},$$

so $Q_{L,\alpha,\bar{\eta}}$ interpolates between the weighted ℓ_1 and weighted ℓ_2 norms of $\boldsymbol{\rho}$. Also, the weights in the norm are *dependent* on the initialization direction $\bar{\eta}$ unless $L = 2$ and $\alpha \rightarrow 0$. In general, $Q_{L,\alpha,\bar{\eta}}$ interpolates the standard ℓ_1 and ℓ_2 norms only if $|[\bar{\eta}]_j|$ is the same for all $j \in [m]$. This result is similar to the observations made in Woodworth et al. [30] which considers a diagonal network with a ‘‘differential’’ structure $f(\mathbf{x}; \mathbf{w}_+, \mathbf{w}_-) = \mathbf{x}^T (\mathbf{w}_+^{\odot L} - \mathbf{w}_-^{\odot L})$. In contrast, our results apply to a more general class of networks, without the need to have the differential structure. In Appendix E.4, we state corollaries of Theorem 4 for linear diagonal networks and linear full-length convolutional networks with even data points. There, we also show that deep matrix sensing with commutative sensor matrices [4] is a special case of our setting.

References

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018.
- [2] Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International Conference on Machine Learning*, pages 244–253, 2018.
- [3] Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. In *International Conference on Learning Representations*, 2019.
- [4] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, pages 7413–7424, 2019.
- [5] Peter Bartlett, Dave Helmbold, and Philip Long. Gradient descent with identity initialization efficiently learns positive definite linear transformations by deep residual networks. In *International Conference on Machine Learning*, pages 521–530, 2018.
- [6] Lenaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *arXiv preprint arXiv:2002.04486*, 2020.
- [7] Lenaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2937–2947, 2019.
- [8] Simon S Du and Wei Hu. Width provably matters in optimization for deep linear neural networks. *arXiv preprint arXiv:1901.08572*, 2019.
- [9] Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*, 2018.
- [10] Simon S Du, Xiyu Zhai, Barnabas Póczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- [11] Daniel Gissin, Shai Shalev-Shwartz, and Amit Daniely. The implicit bias of depth: How incremental learning drives generalization. In *International Conference on Learning Representations*, 2020.
- [12] Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pages 6151–6159, 2017.
- [13] Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, pages 1832–1841, 2018.
- [14] Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems*, pages 9461–9471, 2018.

- [15] Christopher J Hillar and Lek-Heng Lim. Most tensor problems are NP-hard. *Journal of the ACM (JACM)*, 60(6):1–39, 2013.
- [16] Wei Hu, Lechao Xiao, and Jeffrey Pennington. Provable benefit of orthogonal initialization in optimizing deep linear networks. In *International Conference on Learning Representations*, 2020.
- [17] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- [18] Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019.
- [19] Ziwei Ji and Matus Telgarsky. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory*, pages 1772–1798, 2019.
- [20] Ziwei Ji and Matus Telgarsky. A refined primal-dual analysis of the implicit bias. *arXiv preprint arXiv:1906.04540*, 2019.
- [21] Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. *arXiv preprint arXiv:2006.06657*, 2020.
- [22] Lek-Heng Lim. Singular values and eigenvalues of tensors: a variational approach. In *1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, 2005.*, pages 129–132. IEEE, 2005.
- [23] Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020.
- [24] Edward Moroshko, Suriya Gunasekar, Blake Woodworth, Jason D Lee, Nathan Srebro, and Daniel Soudry. Implicit bias in deep linear classification: Initialization scale vs training accuracy. *arXiv preprint arXiv:2007.06738*, 2020.
- [25] 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 *International Conference on Machine Learning*, pages 4683–4692, 2019.
- [26] Mor Shpigel Nacson, Jason Lee, Suriya Gunasekar, Pedro Henrique Pamplona Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3420–3428. PMLR, 2019.
- [27] Mor Shpigel Nacson, Nathan Srebro, and Daniel Soudry. Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3051–3059. PMLR, 2019.
- [28] Samet Oymak and Mahdi Soltanolkotabi. Towards moderate overparameterization: global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 2020.

- [29] 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.
- [30] 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*, 2020.
- [31] Lei Wu, Qingcan Wang, and Chao Ma. Global convergence of gradient descent for deep linear residual networks. In *Advances in Neural Information Processing Systems*, pages 13389–13398, 2019.
- [32] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Michael C Mozer, and Yoram Singer. Identity crisis: Memorization and generalization under extreme overparameterization. *arXiv preprint arXiv:1902.04698*, 2019.
- [33] Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes over-parameterized deep ReLU networks. *arXiv preprint arXiv:1811.08888*, 2018.

Appendix A. Organization and additional notation

Due to page limits, we deferred some parts of the paper to the appendix. We first summarize the organization of the supplementary material and also introduce some additional notation that will be used throughout. In Appendix B, we summarize existing results in the implicit bias/regularization literature that are most relevant to our paper. Appendix C is devoted to showing specific examples of our tensor formulation (1): we show that fully-connected, diagonal, and circular convolutional networks are special cases of our formulation. Next, in Appendix D, we state theorems that were omitted in the main text, which study an extreme overparametrization setting where there is only one data point. Although having $n = 1$ is restrictive, our theorems hold for *any* 2-layer linear tensor networks, including convolutional networks with filter size less than input dimension d , which are not covered by previous results. In Appendix E, we state results on specific network architectures such as fully-connected, diagonal, and convolutional networks. Our simple experiments are presented in Appendix F, and the rest of the sections are the proofs of the theorems and corollaries.

We now summarize some additional notation used in the appendix. Given a matrix \mathbf{A} , we use $\text{vec}(\mathbf{A})$ to denote its vectorization, i.e., the concatenation of all columns of \mathbf{A} . In element indexing of tensors/matrices/vectors, we use \cdot to denote all indices in the corresponding dimension, and $a : b$ to denote all indices from a to b . For example, for a matrix \mathbf{A} , $[\mathbf{A}]_{\cdot,4:6}$ denotes a submatrix that consists of 4th–6th columns of \mathbf{A} . The square bracket notation for indexing overloads with $[a]$ when $a \in \mathbb{N}$, but they will be distinguishable from the context. Since element indices start from 1, we re-define the modulo operation $a \bmod d := a - \lfloor \frac{a-1}{d} \rfloor d \in [d]$ for $a > 0$.

Appendix B. Related works

Gradient flow/descent in separable classification. For linear models $f(\mathbf{x}; \mathbf{z}) = \mathbf{x}^T \mathbf{z}$ with separable data, Soudry et al. [29] show that the GD run on \mathcal{L} drives $\|\mathbf{z}\|$ to ∞ , but \mathbf{z} converges in direction to the ℓ_2 max-margin classifier. The limit direction of \mathbf{z} is aligned with the solution of

$$\text{minimize}_{\mathbf{z} \in \mathbb{R}^d} \quad \|\mathbf{z}\| \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{z} \geq 1 \text{ for } i \in [n], \quad (6)$$

where the norm in the cost is the ℓ_2 norm. Gunasekar et al. [13], Ji and Telgarsky [19, 20], Nacson et al. [26, 27] extend these results to other (stochastic) algorithms and non-separable settings.

Gunasekar et al. [14] study the same problem on linear neural networks and show that GD exhibits different implicit bias depending on the architecture. The authors show that the linear coefficients of the network converges in direction to the solution of (6) with different norms: ℓ_2 norm for linear fully-connected networks, $\ell_{2/L}$ (quasi-)norm for diagonal networks, and DFT-domain $\ell_{2/L}$ (quasi-)norm for convolutional networks with full-length filters. Here, L denotes the depth. We note that [14] assume that GD globally minimizes the loss, and the network parameters and the gradient with respect to the linear coefficients converge in direction. Subsequent results [18, 21] remove such assumptions for linear fully-connected networks.

A recent line of results [21, 23, 25] study general homogeneous models and show divergence of parameters to infinity, monotone increase of smoothed margin, directional convergence and alignment of parameters (see Section 3 for details). Lyu and Li [23] also characterize the limit direction of parameters as the KKT point of a max-margin problem similar to (6), but this characterization does not provide useful insights for the functions $f(\cdot; \Theta)$ represented by specific architectures, because the formulation is in the parameter space Θ . Also, these results require that gradient flow/descent

has already reached 100% training accuracy. Although we study a more restrictive set of networks (i.e., deep linear), we provide a more complete characterization of the implicit bias for the functions $f(\cdot; \Theta)$, without assuming 100% training accuracy.

Gradient flow/descent in linear regression. It is known that for linear models $f(\mathbf{x}; \mathbf{z}) = \mathbf{x}^T \mathbf{z}$, GD converges to the global minimum that is closest in ℓ_2 distance to the initialization (see e.g., [13]). However, relatively less is known for deep networks, even for linear networks. This is partly because the parameters do not diverge to infinity, hence making limit points highly dependent on the initialization; this dependency renders analysis difficult. A related problem of matrix sensing aims to minimize $\sum_{i=1}^n (y_i - \langle \mathbf{A}_i, \mathbf{W}_1 \cdots \mathbf{W}_L \rangle)^2$ over $\mathbf{W}_1, \dots, \mathbf{W}_L \in \mathbb{R}^{d \times d}$. It is shown in [4, 12] that if the sensor matrices \mathbf{A}_i commute and we initialize \mathbf{W}_l to $\alpha \mathbf{I}$, GD finds the minimum nuclear norm solution as $\alpha \rightarrow 0$.

Chizat et al. [7] shows that if a network is zero at initialization, and we scale the network output by a factor of $\alpha \rightarrow \infty$, then the GD dynamics enters a “lazy regime” where the network behaves like a first-order approximation at its initialization, as also seen in results studying kernel approximations of neural networks and convergence of GD in the corresponding RKHS (e.g., [17]).

Woodworth et al. [30] study linear regression with a diagonal network of the form $f(\mathbf{x}; \mathbf{w}_+, \mathbf{w}_-) = \mathbf{x}^T (\mathbf{w}_+^{\odot L} - \mathbf{w}_-^{\odot L})$, where \mathbf{w}_+ and \mathbf{w}_- are identically initialized $\mathbf{w}_+(0) = \mathbf{w}_-(0) = \alpha \bar{\mathbf{w}}$. The authors show that the global minimum reached by GD minimizes a norm-like function which interpolates between (weighted) ℓ_1 norm ($\alpha \rightarrow 0$) and ℓ_2 norm ($\alpha \rightarrow \infty$). In our paper, we consider a more general class of orthogonally decomposable networks, and obtain similar results interpolating between weighted ℓ_1 and ℓ_2 norms. We also remark that our results include the results in [4] as a special case, and we do not assume convergence to global minima, as done in [4, 12, 30].

Appendix C. Tensor representation of neural networks: specific examples

Diagonal networks. An L -layer diagonal network is written as

$$f_{\text{diag}}(\mathbf{x}; \Theta_{\text{diag}}) = \phi(\cdots \phi(\phi(\mathbf{x} \odot \mathbf{w}_1) \odot \mathbf{w}_2) \cdots \odot \mathbf{w}_{L-1})^T \mathbf{w}_L, \quad (7)$$

where $\mathbf{w}_l \in \mathbb{R}^d$ for $l \in [L]$. The representation of f_{diag} as the tensor form (1) is straightforward. Let $\mathbf{M}_{\text{diag}}(\mathbf{x}) \in \mathbb{R}^{d \times \cdots \times d}$ have $[\mathbf{M}_{\text{diag}}(\mathbf{x})]_{j,j,\dots,j} = [\mathbf{x}]_j$, while all the remaining entries of $\mathbf{M}_{\text{diag}}(\mathbf{x})$ are set to zero. We can set $\mathbf{v}_l = \mathbf{w}_l$ for all l , and $\mathbf{M} = \mathbf{M}_{\text{diag}}$ to verify that (1) and (7) are equivalent.

Circular convolutional networks. The tensor formulation (1) includes convolutional networks

$$f_{\text{conv}}(\mathbf{x}; \Theta_{\text{conv}}) = \phi(\cdots \phi(\phi(\mathbf{x} \star \mathbf{w}_1) \star \mathbf{w}_2) \cdots \star \mathbf{w}_{L-1})^T \mathbf{w}_L, \quad (8)$$

where $\mathbf{w}_l \in \mathbb{R}^{k_l}$ with $k_l \leq d$ and $k_L = d$, and \star defines the circular convolution: for any $\mathbf{a} \in \mathbb{R}^d$ and $\mathbf{b} \in \mathbb{R}^k$ ($k \leq d$), we have $\mathbf{a} \star \mathbf{b} \in \mathbb{R}^d$ defined as $[\mathbf{a} \star \mathbf{b}]_i = \sum_{j=1}^k [\mathbf{a}]_{(i+j-1) \bmod d} [\mathbf{b}]_j$, for $i \in [d]$. Define $\mathbf{M}_{\text{conv}}(\mathbf{x}) \in \mathbb{R}^{k_1 \times \cdots \times k_L}$ as $[\mathbf{M}_{\text{conv}}(\mathbf{x})]_{j_1, j_2, \dots, j_L} = [\mathbf{x}]_{(\sum_{l=1}^L j_l - L + 1) \bmod d}$ for $j_l \in [k_l]$, $l \in [L]$. Setting $\mathbf{v}_l = \mathbf{w}_l$ and $\mathbf{M} = \mathbf{M}_{\text{conv}}$, we can verify that (1) and (8) are identical.

Fully-connected networks. An L -layer fully-connected network is defined as

$$f_{\text{fc}}(\mathbf{x}; \Theta_{\text{fc}}) = \phi(\cdots \phi(\phi(\mathbf{x}^T \mathbf{W}_1) \mathbf{W}_2) \cdots \mathbf{W}_{L-1}) \mathbf{w}_L, \quad (9)$$

where $\mathbf{W}_l \in \mathbb{R}^{d_l \times d_{l+1}}$ for $l \in [L-1]$ (we use $d_1 = d$) and $\mathbf{w}_L \in \mathbb{R}^{d_L}$. One can represent f_{fc} as the tensor form (1) by defining parameters $\mathbf{v}_l = \text{vec}(\mathbf{W}_l)$ for $l \in [L-1]$ and $\mathbf{v}_L = \mathbf{w}_L$, and

constructing the tensor $\mathbf{M}_{\text{fc}}(\mathbf{x})$ by a recursive ‘‘block diagonal’’ manner. For example, if $L = 2$, we can define $\mathbf{M}_{\text{fc}}(\mathbf{x}) \in \mathbb{R}^{d_1 d_2 \times d_2}$ to be the Kronecker product of \mathbf{I}_{d_2} and \mathbf{x} .

For deeper networks, we start with $\mathbf{T}_1(\mathbf{x}) := \mathbf{x} \in \mathbb{R}^{d_1}$. Next, define a block diagonal matrix $\mathbf{T}_2(\mathbf{x}) \in \mathbb{R}^{d_1 d_2 \times d_2}$ where the ‘‘diagonals’’ $[\mathbf{T}_2(\mathbf{x})]_{d_1(j-1)+1:d_1 j, j} = \mathbf{T}_1(\mathbf{x})$ for $j \in [d_2]$, while all the other entries are filled with 0. We continue this ‘‘block diagonal’’ procedure, as the following. Having defined $\mathbf{T}_{l-1}(\mathbf{x}) \in \mathbb{R}^{d_1 d_2 \times \dots \times d_{l-2} d_{l-1} \times d_{l-1}}$,

1. Define $\mathbf{T}_l(\mathbf{x}) \in \mathbb{R}^{d_1 d_2 \times \dots \times d_{l-1} d_l \times d_l}$.
2. Set $[\mathbf{T}_l(\mathbf{x})]_{\dots, d_{l-1}(j-1)+1:d_{l-1} j, j} = \mathbf{T}_{l-1}(\mathbf{x}), \forall j \in [d_l]$.
3. Set all the remaining entries of $\mathbf{T}_l(\mathbf{x})$ to zero.

We repeat this process for $l = 2, \dots, L$, and set $\mathbf{M}_{\text{fc}}(\mathbf{x}) := \mathbf{T}_L(\mathbf{x})$. By defining the parameters of the tensor formulation $\mathbf{v}_l = \text{vec}(\mathbf{W}_l)$ for $l \in [L-1]$ and $\mathbf{v}_L = \mathbf{w}_L$, and using the tensor $\mathbf{M}(\mathbf{x}) = \mathbf{M}_{\text{fc}}(\mathbf{x})$, we can check the equivalence of (1) and (9).

Appendix D. Theorems on extremely overparametrized settings

Other than Assumption 1, there is another setting where we can prove a full characterization of limit directions or limit points: when there is one data point ($n = 1$) and the network is 2-layer ($L = 2$). This ‘‘extremely overparametrized’’ case is also motivated by an empirical paper [32] which studies generalization performance of different architectures when there is only one training data point. Here, we present theorems on classification/regression setups. Although having $n = 1$ is restrictive, our results hold for any linear tensor networks including the ones that are not covered by the existing results, hence provide useful insights.

D.1. Classification setting

We first state the theorem for the classification task:

Theorem 5 *Suppose we have a 2-layer linear tensor network (2) and a single data point (\mathbf{x}, y) . Consider the compact SVD $\mathbf{M}(\mathbf{x}) = \mathbf{U}_1 \text{diag}(\mathbf{s}) \mathbf{U}_2^T$, where $\mathbf{U}_1 \in \mathbb{R}^{k_1 \times m}$, $\mathbf{U}_2 \in \mathbb{R}^{k_2 \times m}$, and $\mathbf{s} \in \mathbb{R}^m$ for $m \leq \min\{k_1, k_2\}$. Let $\boldsymbol{\rho}^\infty \in \mathbb{R}^m$ be a solution of the following optimization problem*

$$\text{minimize}_{\boldsymbol{\rho} \in \mathbb{R}^m} \quad \|\boldsymbol{\rho}\|_1 \quad \text{subject to} \quad \mathbf{y} \mathbf{s}^T \boldsymbol{\rho} \geq 1.$$

Assume that there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{v}}_1, \bar{\mathbf{v}}_2$ of the network parameters satisfy $[\mathbf{U}_1^T \bar{\mathbf{v}}_1]_j^2 - [\mathbf{U}_2^T \bar{\mathbf{v}}_2]_j^2 \geq \lambda$ for all $j \in [m]$. Then, \mathbf{v}_1 and \mathbf{v}_2 converge in direction to $\mathbf{U}_1 \boldsymbol{\eta}_1^\infty$ and $\mathbf{U}_2 \boldsymbol{\eta}_2^\infty$, where $|\boldsymbol{\eta}_1^\infty| = |\boldsymbol{\eta}_2^\infty| = |\boldsymbol{\rho}^\infty|^{\odot 1/2}$, and $\text{sign}(\boldsymbol{\eta}_1^\infty) = \text{sign}(y) \odot \text{sign}(\boldsymbol{\eta}_2^\infty)$.

The proof of Theorem 5 can be found in Appendix J. Since $\boldsymbol{\rho}^\infty$ is the minimum ℓ_1 norm solution in the singular value space, the parameters \mathbf{v}_1 and \mathbf{v}_2 converge in direction to the top singular vectors. We would like to emphasize that this theorem can be applied to *any* network architecture that can be represented as a linear tensor network. Recall that the previous result [14] only considers full-length filters ($k_1 = d$), hence providing limited insights on networks with small filters, e.g., $k_1 = 2$. In light of this, we present a corollary in Appendix E.5 showing that linear coefficients of convolutional networks converge in direction to a ‘‘filtered’’ version of \mathbf{x} .

D.2. Regression setting

Next, we present the regression counterpart of Theorem 5, for 2-layer linear tensor networks with a single data point. For this extremely overparametrized setup, we can fully characterize the limit points as functions of initialization $\mathbf{v}_1(0) = \alpha \bar{\mathbf{v}}_1$ and $\mathbf{v}_2(0) = \alpha \bar{\mathbf{v}}_2$.

Theorem 6 *Suppose we have a 2-layer linear tensor network (2) and a single data point (\mathbf{x}, y) . Consider the compact SVD $\mathbf{M}(\mathbf{x}) = \mathbf{U}_1 \text{diag}(\mathbf{s}) \mathbf{U}_2^T$, where $\mathbf{U}_1 \in \mathbb{R}^{k_1 \times m}$, $\mathbf{U}_2 \in \mathbb{R}^{k_2 \times m}$, and $\mathbf{s} \in \mathbb{R}^m$ for $m \leq \min\{k_1, k_2\}$. Assume that there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{v}}_1, \bar{\mathbf{v}}_2$ of the network parameters satisfy $[\mathbf{U}_1^T \bar{\mathbf{v}}_1]_j^2 - [\mathbf{U}_2^T \bar{\mathbf{v}}_2]_j^2 \geq \lambda$ for all $j \in [m]$. Then, gradient flow converges to a global minimizer of the loss \mathcal{L} , and $\mathbf{v}_1(t)$ and $\mathbf{v}_2(t)$ converge to the limit points:*

$$\begin{aligned} \mathbf{v}_1^\infty &= \alpha \mathbf{U}_1 \left(\mathbf{U}_1^T \bar{\mathbf{v}}_1 \odot \cosh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) + \mathbf{U}_2^T \bar{\mathbf{v}}_2 \odot \sinh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) \right) + \alpha (\mathbf{I}_{k_1} - \mathbf{U}_1 \mathbf{U}_1^T) \bar{\mathbf{v}}_1, \\ \mathbf{v}_2^\infty &= \alpha \mathbf{U}_2 \left(\mathbf{U}_1^T \bar{\mathbf{v}}_1 \odot \sinh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) + \mathbf{U}_2^T \bar{\mathbf{v}}_2 \odot \cosh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) \right) + \alpha (\mathbf{I}_{k_2} - \mathbf{U}_2 \mathbf{U}_2^T) \bar{\mathbf{v}}_2, \end{aligned}$$

where g^{-1} is the inverse of the following strictly increasing function

$$g(\nu) = \sum_{j=1}^m [\mathbf{s}]_j \left(\frac{[\mathbf{U}_1^T \bar{\mathbf{v}}_1]_j^2 + [\mathbf{U}_2^T \bar{\mathbf{v}}_2]_j^2}{2} \sinh(2[\mathbf{s}]_j \nu) + [\mathbf{U}_1^T \bar{\mathbf{v}}_1]_j [\mathbf{U}_2^T \bar{\mathbf{v}}_2]_j \cosh(2[\mathbf{s}]_j \nu) \right).$$

The proof can be found in Appendix K. We can observe that as $\alpha \rightarrow 0$, we have $g^{-1} \left(\frac{y}{\alpha^2} \right) \rightarrow \infty$, which results in exponentially faster growth of the $\sinh(\cdot)$ and $\cosh(\cdot)$ for the top singular values. As a result, the top singular vectors dominate the limit points \mathbf{v}_1^∞ and \mathbf{v}_2^∞ as $\alpha \rightarrow 0$, and they do not depend on the initial directions $\bar{\mathbf{v}}_1, \bar{\mathbf{v}}_2$. Experiment results in Section F support this observation.

Appendix E. Results on specific network architectures

In this section, we present corollaries obtained by specializing the theorems in the main text to specific network architectures, as well as a theorem on deep fully-connected networks. We briefly review the linear neural network architectures studied in this section.

Linear fully-connected networks. An L -layer linear fully-connected network is defined as

$$f_{\text{fc}}(\mathbf{x}; \Theta_{\text{fc}}) = \mathbf{x}^T \mathbf{W}_1 \cdots \mathbf{W}_{L-1} \mathbf{w}_L, \quad (10)$$

where $\mathbf{W}_l \in \mathbb{R}^{d_l \times d_{l+1}}$ for $l \in [L-1]$ (we use $d_1 = d$) and $\mathbf{w}_L \in \mathbb{R}^{d_L}$.

Linear diagonal networks. An L -layer linear diagonal network is written as

$$f_{\text{diag}}(\mathbf{x}; \Theta_{\text{diag}}) = (\mathbf{x} \odot \mathbf{w}_1 \odot \cdots \odot \mathbf{w}_{L-1})^T \mathbf{w}_L, \quad (11)$$

where $\mathbf{w}_l \in \mathbb{R}^d$ for $l \in [L]$.

Linear (circular) convolutional networks. An L -layer linear convolutional network is written as

$$f_{\text{conv}}(\mathbf{x}; \Theta_{\text{conv}}) = (\cdots ((\mathbf{x} \star \mathbf{w}_1) \star \mathbf{w}_2) \cdots \star \mathbf{w}_{L-1})^T \mathbf{w}_L, \quad (12)$$

where $\mathbf{w}_l \in \mathbb{R}^{k_l}$ with $k_l \leq d$ and $k_L = d$, and \star defines the circular convolution: for any $\mathbf{a} \in \mathbb{R}^d$ and $\mathbf{b} \in \mathbb{R}^k$ ($k \leq d$), we have $\mathbf{a} \star \mathbf{b} \in \mathbb{R}^d$ defined as $[\mathbf{a} \star \mathbf{b}]_i = \sum_{j=1}^k [\mathbf{a}]_{(i+j-1) \bmod d} [\mathbf{b}]_j$, for $i \in [d]$. In case of $k_l = d$ for all $l \in [L]$, we refer to this network as *full-length* convolutional networks.

Deep matrix sensing. The deep matrix sensing problem considered in Arora et al. [4], Gunasekar et al. [12] aims to minimize the following problem

$$\underset{\mathbf{W}_1, \dots, \mathbf{W}_L \in \mathbb{R}^{d \times d}}{\text{minimize}} \quad \mathcal{L}_{\text{ms}}(\mathbf{W}_1 \cdots \mathbf{W}_L) := \sum_{i=1}^n (y_i - \langle \mathbf{A}_i, \mathbf{W}_1 \cdots \mathbf{W}_L \rangle)^2, \quad (13)$$

where the sensor matrices $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathbb{R}^{d \times d}$ are symmetric. Following Arora et al. [4], Gunasekar et al. [12], we consider sensor matrices $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathbb{R}^{d \times d}$ that commute. To make the problem underdetermined, we assume that $n \leq d$, and \mathbf{A}_i 's are linearly independent.

E.1. Deep fully-connected networks: the $\alpha \rightarrow 0$ limit

Before we state the corollaries, we present another theorem which proves that the linear coefficients $\beta_{\text{fc}}(\Theta_{\text{fc}})$ of deep linear fully-connected networks converge to the minimum ℓ_2 norm solution as $\alpha \rightarrow 0$. We assume for simplicity that $d_1 = d_2 = \dots = d_L = d$ in this section, but we can extend it for $d_l \geq d$ without too much difficulty. Recall $f_{\text{fc}}(\mathbf{x}; \Theta_{\text{fc}}) = \mathbf{x}^T \mathbf{W}_1 \cdots \mathbf{W}_{L-1} \mathbf{w}_L$. We minimize the training loss \mathcal{L} with initialization $\mathbf{W}_l(0) = \alpha \bar{\mathbf{W}}_l$ for $l \in [L-1]$ and $\mathbf{w}_L(0) = \alpha \bar{\mathbf{w}}_L$.

Theorem 7 *Assume that initial directions $\bar{\mathbf{W}}_1, \dots, \bar{\mathbf{W}}_{L-1}, \bar{\mathbf{w}}_L$ satisfy (1) $\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l \succeq \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T$ for $l \in [L-2]$, and (2) there exists $\lambda > 0$ such that $\bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T \succeq \lambda \mathbf{I}_d$. Then, the gradient flow converges to a global minimum, and $\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} \beta_{\text{fc}}(\Theta_{\text{fc}}(t)) = \mathbf{X}^T (\mathbf{X} \mathbf{X}^T)^{-1} \mathbf{y}$.*

The proof is presented in Appendix L. Theorem 7 shows that in the limit $\alpha \rightarrow 0$, linear fully-connected networks have bias towards the minimum ℓ_2 norm solution, regardless of the depth. This is consistent with the results shown for classification. We also note that the convergence to a global minimum holds regardless of $\alpha > 0$, and our sufficient conditions on global convergence generalize the zero-asymmetric initialization scheme proposed in Wu et al. [31].

E.2. Corollary of Theorem 1

Corollary 8 *Consider an L -layer linear fully-connected network (10). If the training loss satisfies $\mathcal{L}(\Theta_{\text{fc}}(t_0)) < 1$ for some $t_0 \geq 0$, then $\beta_{\text{fc}}(\Theta_{\text{fc}}(t))$ converges in a direction that aligns with the solution of the following optimization problem*

$$\underset{\mathbf{z} \in \mathbb{R}^d}{\text{minimize}} \quad \|\mathbf{z}\|_2^2 \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{z} \geq 1, \quad \forall i \in [n].$$

Corollary 8 shows that whenever the network separates the data correctly, the linear coefficients $\beta_{\text{fc}}(\Theta_{\text{fc}})$ convergence in direction to the ℓ_2 max-margin classifier. Note that this corollary does not require the directional convergence of $\mathbf{X}^T \mathbf{r}$, which is different from Theorem 1. In fact, this corollary also appears in Ji and Telgarsky [21], but we provide an alternative proof based on our tensor formulation. The proof of Corollary 8 can be found in Appendix G.

E.3. Corollaries of Theorem 2

Corollary 9 *Consider an L -layer linear diagonal network (11). If there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{w}}_1, \dots, \bar{\mathbf{w}}_L$ of the network parameters satisfy $[\bar{\mathbf{w}}_l]_j^2 - [\bar{\mathbf{w}}_L]_j^2 \geq \lambda$ for all $l \in [L-1]$ and $j \in [d]$, then $\beta_{\text{diag}}(\Theta_{\text{diag}}(t))$ converges in a direction that aligns with a stationary point \mathbf{z}^∞ of*

$$\underset{\mathbf{z} \in \mathbb{R}^d}{\text{minimize}} \quad \|\mathbf{z}\|_{2/L} \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{z} \geq 1, \quad \forall i \in [n].$$

For full-length convolutional networks, we define $\mathbf{F} \in \mathbb{C}^{d \times d}$ to be the matrix of discrete Fourier transform basis $[\mathbf{F}]_{j,k} = \frac{1}{\sqrt{d}} \exp(-\frac{\sqrt{-1} \cdot 2\pi(j-1)(k-1)}{d})$. Note that $\mathbf{F}^* = \mathbf{F}^{-1}$, and both \mathbf{F} and \mathbf{F}^* are symmetric, but not Hermitian.

Corollary 10 *Consider an L -layer linear full-length convolutional network (12). If there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{w}}_1, \dots, \bar{\mathbf{w}}_L$ of the network parameters satisfy $|[\mathbf{F}\bar{\mathbf{w}}_l]_j|^2 - |[\mathbf{F}\bar{\mathbf{w}}_L]_j|^2 \geq \lambda$ for all $l \in [L-1]$ and $j \in [d]$, then $\beta_{\text{conv}}(\Theta_{\text{conv}}(t))$ converges in a direction that aligns with a stationary point \mathbf{z}^∞ of*

$$\text{minimize}_{\mathbf{z} \in \mathbb{R}^d} \quad \|\mathbf{F}\mathbf{z}\|_{2/L} \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{z} \geq 1, \quad \forall i \in [n].$$

Corollary 9 shows that in the limit, linear diagonal network finds a sparse solution \mathbf{z} that is a stationary point of the $\ell_{2/L}$ max-margin classification problem. Corollary 10 has a similar conclusion except that the standard $\ell_{2/L}$ norm is replaced with DFT-domain $\ell_{2/L}$ norm. By specifying mild conditions on initialization (see discussion below Theorem 2), these corollaries remove the convergence assumptions required in Gunasekar et al. [14]. The proofs of Corollaries 9 and 10 are in Appendix H.

E.4. Corollaries of Theorem 4

In this subsection, we apply Theorem 4 to linear diagonal networks, linear full-length convolutional networks with even data, and deep matrix sensing. The proofs of the corollaries can be found in Appendix I.

Corollary 11 *Consider an L -layer linear diagonal network (11). For some $\lambda > 0$, choose any vector $\bar{\mathbf{w}} \in \mathbb{R}^d$ satisfying $[\bar{\mathbf{w}}]_j^2 \geq \lambda$ for all $j \in [d]$, and choose initial directions $\bar{\mathbf{w}}_l = \bar{\mathbf{w}}$ for $l \in [L-1]$ and $\bar{\mathbf{w}}_L = \mathbf{0}$. Then, the linear coefficients $\beta_{\text{diag}}(\Theta_{\text{diag}}(t))$ converge to the solution \mathbf{z}^∞ of*

$$\text{minimize}_{\mathbf{z} \in \mathbb{R}^d} \quad Q_{L,\alpha,\bar{\mathbf{w}}}(\mathbf{z}) := \alpha^2 \sum_{j=1}^d [\bar{\mathbf{w}}]_j^2 H_L \left(\frac{[z]_j}{\alpha^L [\bar{\mathbf{w}}]_j^L} \right) \quad \text{subject to} \quad \mathbf{X}\mathbf{z} = \mathbf{y}.$$

Recall that the original statement of Assumption 1 allows the matrices $\mathbf{S}, \mathbf{U}_1, \dots, \mathbf{U}_L$ to be complex, but Theorem 4 poses another assumption that these matrices are real. In applying Theorem 2 to convolutional networks to get Corollary 10, we used the fact that the data tensor $\mathbf{M}_{\text{conv}}(\mathbf{x})$ of a linear full-length convolutional network satisfies Assumption 1 with $\mathbf{S} = d^{\frac{L-1}{2}} \mathbf{F}$ and $\mathbf{U}_1 = \dots = \mathbf{U}_L = \mathbf{F}^*$, where $\mathbf{F} \in \mathbb{C}^{d \times d}$ is the matrix of discrete Fourier transform basis $[\mathbf{F}]_{j,k} = \frac{1}{\sqrt{d}} \exp(-\frac{\sqrt{-1} \cdot 2\pi(j-1)(k-1)}{d})$ and \mathbf{F}^* is the complex conjugate of \mathbf{F} . Note that these are complex matrices, so one cannot directly apply Theorem 4 to convolutional networks. However, it turns out that if the data and initialization are even, we can derive a corollary for convolutional networks.

We say that a vector is *even* when it satisfies the even symmetry, as in even functions. More concretely, a vector $\mathbf{x} \in \mathbb{R}^d$ is even if $[\mathbf{x}]_{j+2} = [\mathbf{x}]_{d-j}$ for $j = 0, \dots, \lfloor \frac{d-3}{2} \rfloor$; i.e., the vector has the even symmetry around its ‘‘origin’’ $[\mathbf{x}]_1$. From the definition of the matrix $\mathbf{F} \in \mathbb{C}^{d \times d}$, it is straightforward to check that if \mathbf{x} is real and even, then its DFT $\mathbf{F}\mathbf{x}$ is also real and even (see Appendix I.4 for details).

Corollary 12 Consider an L -layer linear full-length convolutional network (12). Assume that the data points $\{\mathbf{x}_i\}_{i=1}^n$ are all even. For some $\lambda > 0$, choose any even vector $\bar{\mathbf{w}}$ satisfying $[\mathbf{F}\bar{\mathbf{w}}]_j^2 \geq \lambda$ for all $j \in [d]$, and choose initial directions $\bar{\mathbf{w}}_l = \bar{\mathbf{w}}$ for $l \in [L-1]$ and $\bar{\mathbf{w}}_L = \mathbf{0}$. Then, the linear coefficients $\beta_{\text{conv}}(\Theta_{\text{conv}}(t))$ converge to the solution \mathbf{z}^∞ of

$$\underset{\mathbf{z} \in \mathbb{R}^d, \text{ even}}{\text{minimize}} \quad Q_{L,\alpha,\mathbf{F}\bar{\mathbf{w}}}(\mathbf{F}\mathbf{z}) := \alpha^2 \sum_{j=1}^d [\mathbf{F}\bar{\mathbf{w}}]_j^2 H_L \left(\frac{[\mathbf{F}\mathbf{z}]_j}{\alpha^L |[\mathbf{F}\bar{\mathbf{w}}]_j|^L} \right) \quad \text{subject to} \quad \mathbf{X}\mathbf{z} = \mathbf{y}.$$

Corollaries 11 and 12 show that the interpolation between minimum weighted ℓ_1 and weighted ℓ_2 solutions occurs for diagonal networks, and also for convolutional networks (in DFT domain, with the restriction of even symmetry). The conclusion of Corollary 11 is similar to the results in Woodworth et al. [30], but the network architecture (11) considered in our corollary is a slightly different from the ‘‘differential’’ network $f(\mathbf{x}; \mathbf{w}_+, \mathbf{w}_-) = \mathbf{x}^T (\mathbf{w}_+^{\odot L} - \mathbf{w}_-^{\odot L})$ in Woodworth et al. [30].

As mentioned in the main text, we can actually show that the matrix sensing result in Arora et al. [4] is a special case of our Theorem 4. Given any symmetric matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$, let $\text{eig}(\mathbf{M}) \in \mathbb{R}^d$ be the d -dimensional vector of eigenvalues of \mathbf{M} .

Corollary 13 Consider the depth- L deep matrix sensing problem (13). Let \mathbf{A}_i ’s be symmetric, and assume $\mathbf{A}_1, \dots, \mathbf{A}_n$ commute. For $\alpha > 0$, choose initialization $\mathbf{W}_l(0) = \alpha \mathbf{I}_d$ for $l \in [L-1]$ and $\mathbf{W}_L(0) = \mathbf{0}$. Then, the product $\mathbf{W}_L(t) \cdots \mathbf{W}_1(t)$ converge to the solution \mathbf{M}^∞ of

$$\underset{\mathbf{M} \in \mathbb{R}^{d \times d}, \text{ symmetric}}{\text{minimize}} \quad Q_{L,\alpha}(\text{eig}(\mathbf{M})) := \alpha^2 \sum_{j=1}^d H_L \left(\frac{[\text{eig}(\mathbf{M})]_j}{\alpha^L} \right) \quad \text{subject to} \quad \mathcal{L}_{\text{ms}}(\mathbf{M}) = 0.$$

Under an additional assumption that \mathbf{A}_i ’s are positive semidefinite, Theorem 2 in Arora et al. [4] studies the initialization $\mathbf{W}_l(0) = \alpha \mathbf{I}_d$ for all $l \in [L]$, and shows that the limit point of $\mathbf{W}_L \cdots \mathbf{W}_1$ converges to the minimum nuclear norm solution as $\alpha \rightarrow 0$. We remove the assumption of positive definiteness of \mathbf{A}_i ’s and let $\mathbf{W}_L(0) = \mathbf{0}$, to show a complete characterization of the solution found by gradient flow, which interpolates between the minimum nuclear norm (i.e., Schatten 1-norm) solution (when $\alpha \rightarrow 0$) and the minimum Frobenius norm (i.e., Schatten 2-norm) solution (when $\alpha \rightarrow \infty$).

E.5. Corollary of Theorem 5

Recall that Theorem 5 can be applied to any 2-layer networks that can be represented as linear tensor networks. Examples include the convolutional networks that are not full-length (i.e., filter size $k_1 < d$), which are not covered by the previous result [14]. Here, we present the characterization of convergence directions of $\beta_{\text{conv}}(\Theta_{\text{conv}}(t))$ for small-filter cases: $k_1 = 1$ and $k_1 = 2$.

Corollary 14 Consider a 2-layer linear convolutional network (12) with $k_1 = 1$ and a single data point (\mathbf{x}, y) . If there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{w}}_1$ and $\bar{\mathbf{w}}_2$ of the network parameters satisfy $\|\mathbf{x}\|^2 \bar{\mathbf{v}}_1^2 - (\mathbf{x}^T \bar{\mathbf{v}}_2)^2 \geq \|\mathbf{x}\|^2 \lambda$, then $\beta_{\text{conv}}(\Theta_{\text{conv}}(t))$ converges in direction that aligns with $y\mathbf{x}$.

Consider a 2-layer linear convolutional network (12) with $k_1 = 2$ and a single data point (\mathbf{x}, y) . Let $\overleftarrow{\mathbf{x}} := [[\mathbf{x}]_2 \cdots [\mathbf{x}]_d \ [\mathbf{x}]_1]$, and $\overrightarrow{\mathbf{x}} := [[\mathbf{x}]_d \ [\mathbf{x}]_1 \cdots [\mathbf{x}]_{d-1}]$. If there exists $\lambda > 0$ such that the initial directions $\bar{\mathbf{w}}_1$ and $\bar{\mathbf{w}}_2$ of the network parameters satisfy

$$([\bar{\mathbf{v}}_1]_1 + [\bar{\mathbf{v}}_1]_2)^2 - \frac{((\mathbf{x} + \overleftarrow{\mathbf{x}})^T \bar{\mathbf{v}}_2)^2}{\|\mathbf{x}\|_2^2 + \mathbf{x}^T \overleftarrow{\mathbf{x}}} \geq \lambda, \quad \text{and} \quad ([\bar{\mathbf{v}}_1]_1 - [\bar{\mathbf{v}}_1]_2)^2 - \frac{((\mathbf{x} - \overleftarrow{\mathbf{x}})^T \bar{\mathbf{v}}_2)^2}{\|\mathbf{x}\|_2^2 - \mathbf{x}^T \overleftarrow{\mathbf{x}}} \geq \lambda,$$

then $\beta_{\text{conv}}(\Theta_{\text{conv}}(t))$ converges in a direction that aligns with a filtered version of \mathbf{x} :

$$\lim_{t \rightarrow \infty} \frac{\beta_{\text{conv}}(\Theta_{\text{conv}}(t))}{\|\beta_{\text{conv}}(\Theta_{\text{conv}}(t))\|_2} \propto \begin{cases} 2y\mathbf{x} + y\overleftarrow{\mathbf{x}} + y\overrightarrow{\mathbf{x}} & \text{if } \mathbf{x}^T \overleftarrow{\mathbf{x}} > 0, \\ 2y\mathbf{x} - y\overleftarrow{\mathbf{x}} - y\overrightarrow{\mathbf{x}} & \text{if } \mathbf{x}^T \overleftarrow{\mathbf{x}} < 0. \end{cases}$$

Corollary 14 shows that if the filter size is $k_1 = 1$, then the limit direction of $\beta_{\text{conv}}(\Theta_{\text{conv}})$ is always the ℓ_2 max-margin classifier. Note that this is quite different from the case $k_1 = d$ which converges to the DFT-domain ℓ_1 max-margin classifier. However, for $1 < k_1 < d$, it is difficult to characterize the limit direction as the max-margin classifier of some common norms. Rather, the limit directions of $\beta_{\text{conv}}(\Theta_{\text{conv}})$ correspond to a “filtered” version of the data point, and the weights of the filter depend on the data point \mathbf{x} . For $k_1 = 2$, the filter is a low-pass filter if the autocorrelation $\mathbf{x}^T \overleftarrow{\mathbf{x}}$ of \mathbf{x} is positive, and high-pass if the autocorrelation is negative. For $k_1 > 2$, the filter weights are more complicated to characterize in terms of \mathbf{x} , and the filter length increases as k_1 increases. We prove Corollary 14 in Appendix J.

Appendix F. Experiments

Regression. To fully visualize the trajectory of linear coefficients, we run simple experiments with 2-layer linear fully-connected/diagonal/convolutional networks with a single 2-dimensional data point $(\mathbf{x}, y) = ([1 \ 2], 1)$. For this dataset, the minimum ℓ_2 norm solution (corresponding to fully-connected networks) of the regression problem is $[0.2 \ 0.4]$, whereas the minimum ℓ_1 norm solution (corresponding to diagonal) is $[0 \ 0.5]$ and the minimum DFT-domain ℓ_1 norm solution (corresponding to convolutional) is $[0.33 \ 0.33]$. We randomly pick four directions $\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_4 \in \mathbb{R}^2$, and choose initial directions of the network parameters in a way that their linear coefficients at initialization are exactly $\beta(\Theta(0)) = \alpha^2 \bar{\mathbf{z}}_j$. With varying initial scales $\alpha \in \{0.01, 0.5, 1\}$, we run GD with small step size $\eta = 10^{-3}$ for large enough number of iterations $T = 5 \times 10^3$. Figure 1 plot the trajectories of $\beta(\Theta)$ (appropriately clipped for visual clarity) as well as the predicted limit points (Theorem 6). We observe that even though the networks start at the same linear coefficients $\alpha^2 \bar{\mathbf{z}}_j$, they evolve differently due to different architectures. Note that the prediction of limit points is accurate, and the solution found by GD is less dependent on initial directions when α is small.

Classification. It is shown in the existing works as well as in Section 3 that the limit directions of linear coefficients are independent of the initialization. Is this also true in practice? To see this, we run a set of toy experiments on classification with two data points $(\mathbf{x}_1, y_1) = ([1 \ 2], +1)$ and $(\mathbf{x}_2, y_2) = ([0 \ -3], -1)$. One can check that the max-margin classifiers for this problem are in the same directions to the corresponding min-norm solutions in the regression problem above. We use the same networks as in regression, and the same set of initial directions satisfying $\beta(\Theta(0)) = \alpha^2 \bar{\mathbf{z}}_j$. With initial scales $\alpha \in \{0.01, 0.5, 1\}$, we run GD with step size $\eta = 5 \times 10^{-4}$ for $T = 2 \times 10^6$ iterations. All experiments reached $\mathcal{L}(\Theta) \lesssim 10^{-5}$ at the end. The trajectories are plotted in Figure 1. We find that unlike theory, the actual coefficients are quite dependent on initialization, because we do not train the network all the way to zero loss. This observation is also consistent with a recent analysis [24] for diagonal networks, and suggests that understanding the behavior of iterates after a finite number of steps is an important future work.

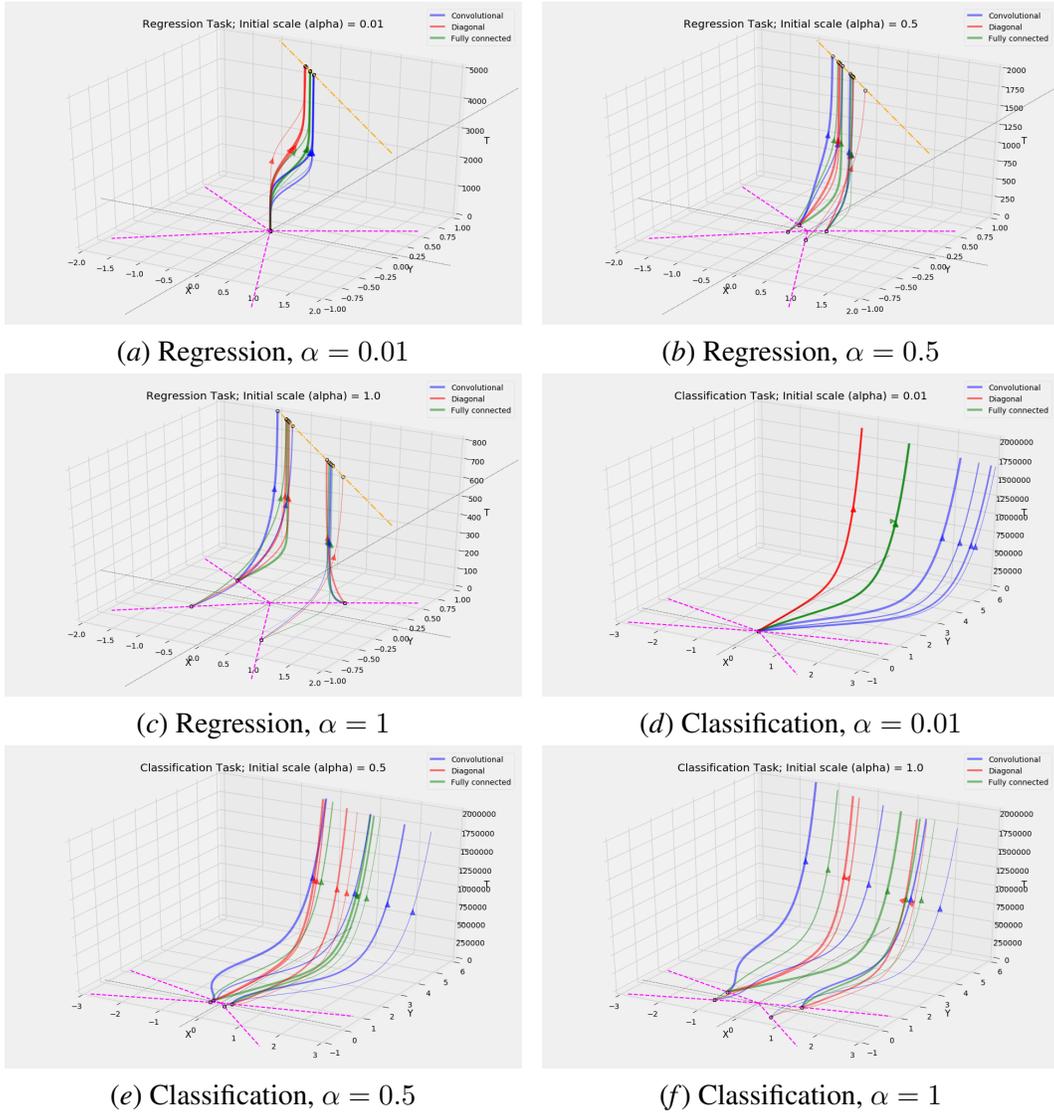


Figure 1: Gradient descent trajectories of linear coefficients of linear fully-connected, diagonal, and convolutional networks on regression/classification tasks, initialized with different initial scales $\alpha = 0.01, 0.5, 1$. Networks are initialized at the same coefficients (circles on purple lines), but follow different trajectories due to implicit biases of networks induced from their architecture. Figures 1a, 1b and 1c show that our theoretical predictions (Theorem 6) on limit points (circles on yellow line, the set of global minima) agree with the solution found by GD.

Appendix G. Proofs of Theorem 1 and Corollary 8

G.1. Proof of Theorem 1

The proof of Theorem 1 is outlined as follows. First, using the directional convergence and alignment results in Ji and Telgarsky [21], we prove that each of our network parameters v_i converges

in direction, and it aligns with its corresponding negative gradient $-\nabla_{\mathbf{v}_l} \mathcal{L}$. Then, we prove that the directions of \mathbf{v}_l 's are actually singular vectors of $\mathbf{M}(-\mathbf{u}^\infty)$, where $\mathbf{u}^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{X}^T \mathbf{r}(t)}{\|\mathbf{X}^T \mathbf{r}(t)\|_2}$.

Since a linear tensor network is an L -homogeneous polynomial of $\mathbf{v}_1, \dots, \mathbf{v}_L$, it satisfies the assumptions required for Theorems 3.1 and 4.1 in Ji and Telgarsky [21]. These theorems imply that if the gradient flow satisfies $\mathcal{L}(\Theta(t_0)) < 1$ for some $t_0 \geq 0$, then $\Theta(t)$ converges in direction, and the direction aligns with $-\nabla_{\Theta} \mathcal{L}(\Theta(t))$; that is,

$$\lim_{t \rightarrow \infty} \|\Theta(t)\|_2 = \infty, \quad \lim_{t \rightarrow \infty} \frac{\Theta(t)}{\|\Theta(t)\|_2} = \Theta^\infty, \quad \lim_{t \rightarrow \infty} \frac{\Theta(t)^T \nabla_{\Theta} \mathcal{L}(\Theta(t))}{\|\Theta(t)\|_2 \|\nabla_{\Theta} \mathcal{L}(\Theta(t))\|_2} = -1. \quad (14)$$

For linear tensor networks (2), the parameter Θ is the concatenation of all parameter vectors $\mathbf{v}_1, \dots, \mathbf{v}_L$, so (14) holds for $\Theta = [\mathbf{v}_1^T \ \dots \ \mathbf{v}_L^T]^T$.

Now, recall that by the definition of the linear tensor network, we have the following gradient flow

$$\dot{\mathbf{v}}_l = \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L).$$

Note that we can apply this to calculate the rate of growth of $\|\mathbf{v}_l\|_2^2$:

$$\begin{aligned} \frac{d}{dt} \|\mathbf{v}_l\|_2^2 &= 2\mathbf{v}_l^T \dot{\mathbf{v}}_l = 2\mathbf{v}_l^T \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L) \\ &= 2\mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{v}_l, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L) \\ &= \frac{d}{dt} \|\mathbf{v}_{l'}\|_2^2 \quad \text{for any } l' \in [L], \end{aligned}$$

so the rate at which $\|\mathbf{v}_l\|_2^2$ grows over time is the same for all layers $l \in [L]$. By the definition of Θ and (14), we have

$$\|\Theta\|_2^2 = \sum_{l=1}^L \|\mathbf{v}_l\|_2^2 \rightarrow \infty,$$

which then implies

$$\lim_{t \rightarrow \infty} \|\mathbf{v}_l(t)\|_2 \rightarrow \infty, \quad \lim_{t \rightarrow \infty} \frac{\|\Theta(t)\|_2}{\|\mathbf{v}_l(t)\|_2} = \sqrt{\frac{\|\Theta(t)\|_2^2}{\|\mathbf{v}_l(t)\|_2^2}} = \sqrt{L},$$

for all $l \in [L]$. Now, let \mathcal{I}_l be the set of indices that correspond to the components of \mathbf{v}_l in Θ . It follows from (14) that

$$\lim_{t \rightarrow \infty} \frac{\mathbf{v}_l(t)}{\|\mathbf{v}_l(t)\|_2} = \lim_{t \rightarrow \infty} \frac{\mathbf{v}_l(t)}{\|\Theta(t)\|_2} \frac{\|\Theta(t)\|_2}{\|\mathbf{v}_l(t)\|_2} = \lim_{t \rightarrow \infty} \frac{[\Theta(t)]_{\mathcal{I}_l}}{\|\Theta(t)\|_2} \frac{\|\Theta(t)\|_2}{\|\mathbf{v}_l(t)\|_2} = \sqrt{L} [\Theta^\infty]_{\mathcal{I}_l},$$

thus showing the directional convergence of \mathbf{v}_l 's.

Next, it follows from directional convergence of Θ and its alignment with $-\nabla_{\Theta} \mathcal{L}(\Theta)$ (14) that $\nabla_{\Theta} \mathcal{L}(\Theta)$ also converges in direction, in the opposite direction of Θ . By comparing the components in \mathcal{I}_l 's, we get that $\nabla_{\mathbf{v}_l} \mathcal{L}(\Theta)$ converges in the opposite direction of \mathbf{v}_l .

For any $l \in [L]$, now let $\mathbf{v}_l^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{v}_l(t)}{\|\mathbf{v}_l(t)\|_2}$. Also recall the assumption that $\mathbf{X}^T \mathbf{r}(t)$ converges in direction; let the unit vector $\mathbf{u}^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{X}^T \mathbf{r}(t)}{\|\mathbf{X}^T \mathbf{r}(t)\|_2}$ be the limit direction. By the gradient flow dynamics of \mathbf{v}_l , we have

$$\mathbf{v}_l^\infty \propto -\nabla_{\mathbf{v}_l} \mathcal{L}(\Theta^\infty) = \mathbf{M}(-\mathbf{u}^\infty) \circ (\mathbf{v}_1^\infty, \dots, \mathbf{v}_{l-1}^\infty, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}^\infty, \dots, \mathbf{v}_L^\infty),$$

for all $l \in [L]$. Note that this equation has the same form as (4), the definition of singular vectors in tensors. So this proves that $(\mathbf{v}_1^\infty, \dots, \mathbf{v}_L^\infty)$ are singular vectors of $\mathbf{M}(-\mathbf{u}^\infty)$.

G.2. Proof of Corollary 8

The proof proceeds as follows. First, we will show using the structure of the data tensor \mathbf{M}_{fc} that the limit direction of linear coefficients $\beta_{\text{fc}}(\Theta_{\text{fc}}^\infty)$ is proportional to $c\mathbf{u}^\infty$, where c is a nonzero scalar and \mathbf{u}^∞ is the limit direction of $\mathbf{X}^T \mathbf{r}$. Then, through a closer look at \mathbf{u}^∞ and c , we will prove that $\beta_{\text{fc}}(\Theta_{\text{fc}}^\infty)$ is in fact a conic combination of the support vectors (i.e., the data points with the minimum margins). Finally, we will compare $\beta_{\text{fc}}(\Theta_{\text{fc}}^\infty)$ with the KKT conditions of the ℓ_2 max-margin classification problem and conclude that $\beta_{\text{fc}}(\Theta_{\text{fc}}^\infty)$ must be in the same direction as the ℓ_2 max-margin classifier.

Due to the way how the data tensor \mathbf{M}_{fc} is constructed for fully-connected networks (Appendix C), we always have

$$-\nabla_{\mathbf{v}_1} \mathcal{L}(\Theta_{\text{fc}}) = \mathbf{M}_{\text{fc}}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{I}_{k_1}, \mathbf{v}_2, \dots, \mathbf{v}_L) \in \text{span} \left\{ \begin{matrix} \left[\begin{matrix} \mathbf{X}^T \mathbf{r} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{matrix} \right], \left[\begin{matrix} \mathbf{0} \\ \mathbf{X}^T \mathbf{r} \\ \vdots \\ \mathbf{0} \end{matrix} \right], \dots, \left[\begin{matrix} \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{X}^T \mathbf{r} \end{matrix} \right] \end{matrix} \right\}.$$

From Theorem 1, we have directional convergence of \mathbf{v}_1 and its alignment with $-\nabla_{\mathbf{v}_1} \mathcal{L}(\Theta_{\text{fc}})$. This means that the limit direction \mathbf{v}_1^∞ , which is a fixed vector, must be also in the span of vectors written above. This implies that $\mathbf{X}^T \mathbf{r}$ must also converge to some direction, say $\mathbf{u}^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{X}^T \mathbf{r}(t)}{\|\mathbf{X}^T \mathbf{r}(t)\|_2}$.

Now recall the definition of \mathbf{v}_1 in case of the fully-connected network: $\mathbf{v}_1 = \text{vec}(\mathbf{W}_1)$. So, by reshaping \mathbf{v}_1^∞ into its original $d_1 \times d_2$ matrix form \mathbf{W}_1^∞ , we have

$$\mathbf{W}_1^\infty \propto \mathbf{u}^\infty \mathbf{q}^T,$$

for some $\mathbf{q} \in \mathbb{R}^{d_2}$. This implies that the linear coefficients $\beta_{\text{fc}}(\Theta_{\text{fc}})$ of the network converge in direction to

$$\beta_{\text{fc}}(\Theta_{\text{fc}}^\infty) = \mathbf{W}_1^\infty \mathbf{W}_2^\infty \dots \mathbf{W}_{L-1}^\infty \mathbf{w}_L^\infty \propto \mathbf{u}^\infty \mathbf{q}^T \mathbf{W}_2^\infty \dots \mathbf{W}_{L-1}^\infty \mathbf{w}_L^\infty = c\mathbf{u}^\infty, \quad (15)$$

where c is some nonzero real number.

Let us now take a closer look at the vector \mathbf{u}^∞ , the limit direction of $\mathbf{X}^T \mathbf{r}$. Recall from Section 2.1 that for any $i \in [n]$,

$$[\mathbf{r}]_i = -y_i \exp(-y_i f_{\text{fc}}(\mathbf{x}_i; \Theta_{\text{fc}})) = -y_i \exp(-y_i \mathbf{x}_i^T \beta_{\text{fc}}(\Theta_{\text{fc}})),$$

in case of classification. Recall that $\|\beta_{\text{fc}}(\Theta_{\text{fc}}(t))\|_2 \rightarrow \infty$ while converging to a certain direction $\beta_{\text{fc}}(\Theta_{\text{fc}}^\infty)$. This means that if

$$y_j \mathbf{x}_j^T \beta_{\text{fc}}(\Theta_{\text{fc}}^\infty) > y_i \mathbf{x}_i^T \beta_{\text{fc}}(\Theta_{\text{fc}}^\infty)$$

for any $i, j \in [n]$, then

$$\lim_{t \rightarrow \infty} \frac{\exp(-y_j \mathbf{x}_j^T \beta_{\text{fc}}(\Theta_{\text{fc}}(t)))}{\exp(-y_i \mathbf{x}_i^T \beta_{\text{fc}}(\Theta_{\text{fc}}(t)))} = 0. \quad (16)$$

Take i to be the index of any support vectors, i.e., any i that attains the minimum $y_i \mathbf{x}_i^T \boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty)$ among all data points. Using such an i , the observation (16) implies that $\lim_{t \rightarrow \infty} [\mathbf{r}(t)]_j = 0$ for any \mathbf{x}_j that is not a support vector. Thus, by the argument above, \mathbf{u}^∞ can in fact be written as

$$\mathbf{u}^\infty \propto \lim_{t \rightarrow \infty} \sum_{i=1}^n \mathbf{x}_i [\mathbf{r}(t)]_i = - \sum_{i=1}^n \nu_i y_i \mathbf{x}_i, \quad (17)$$

where $\nu_i \geq 0$ for all $i \in [n]$, and $\nu_j = 0$ for \mathbf{x}_j 's that are *not* support vectors. Combining (17) and (15),

$$\boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty) \propto -c \sum_{i=1}^n \nu_i y_i \mathbf{x}_i. \quad (18)$$

Recall that we do not yet know whether c , introduced in (15), is positive or negative; we will now show that c has to be negative. From Lyu and Li [23], we know that $\mathcal{L}(\boldsymbol{\Theta}_{\text{fc}}(t)) \rightarrow 0$, which implies that $y_i \mathbf{x}_i^T \boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty) > 0$ for all $i \in [n]$. However, if $c > 0$, then (18) implies that $\boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty)$ is inside a cone \mathcal{K} defined as

$$\mathcal{K} := \left\{ \sum_{i=1}^n \gamma_i y_i \mathbf{x}_i \mid \gamma_i \leq 0, \forall i \in [n] \right\}.$$

Note that the polar cone of \mathcal{K} , denoted as \mathcal{K}° , is

$$\mathcal{K}^\circ := \{ \mathbf{z} \mid \boldsymbol{\beta}^T \mathbf{z} \leq 0, \forall \boldsymbol{\beta} \in \mathcal{K} \} = \{ \mathbf{z} \mid y_i \mathbf{x}_i^T \mathbf{z} \geq 0, \forall i \in [n] \}.$$

It is known that $\mathcal{K} \cap \mathcal{K}^\circ = \{\mathbf{0}\}$ for any convex cone \mathcal{K} and its polar cone \mathcal{K}° . Therefore, having $c > 0$ implies that there exists some $i \in [n]$ such that $y_i \mathbf{x}_i^T \boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty) < 0$, which contradicts the fact that the loss goes to zero as $t \rightarrow \infty$. Therefore, c in (15) and (18) must be negative:

$$\boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty) \propto \sum_{i=1}^n \nu_i y_i \mathbf{x}_i, \quad (19)$$

for $\nu_i \geq 0$ for all $i \in [n]$ and $\nu_j = 0$ for all \mathbf{x}_j 's that are not support vectors.

Finally, compare (19) with the KKT conditions of the following optimization problem:

$$\underset{\mathbf{z}}{\text{minimize}} \quad \|\mathbf{z}\|_2^2 \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{z} \geq 1, \quad \forall i \in [n].$$

The KKT conditions of this problem are

$$\mathbf{z} = \sum_{i=1}^n \mu_i y_i \mathbf{x}_i, \quad \text{and} \quad \mu_i \geq 0, \quad \mu_i (1 - y_i \mathbf{x}_i^T \mathbf{z}) = 0 \quad \text{for all } i \in [n],$$

where μ_1, \dots, μ_n are the dual variables. Note that this is (up to scaling) satisfied by $\boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty)$ (19), if we replace μ_i 's with ν_i 's. This finishes the proof that $\boldsymbol{\beta}_{\text{fc}}(\boldsymbol{\Theta}_{\text{fc}}^\infty)$ is aligned with the ℓ_2 max-margin classifier.

Appendix H. Proofs of Theorem 2 and Corollaries 9 & 10

H.1. Proof of Theorem 2

H.1.1. CONVERGENCE OF LOSS TO ZERO

Since Theorem 2 does not assume the existence of $t_0 \geq 0$ satisfying $\mathcal{L}(\Theta(t_0)) < 1$, we need to first show that given the conditions on initialization, the training loss $\mathcal{L}(\Theta(t))$ converges to zero. Recall from Section 2.1 that

$$\dot{\mathbf{v}}_l = -\nabla_{\mathbf{v}_l} \mathcal{L}(\Theta) = \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L).$$

Applying the structure (5) in Assumption 1, we get

$$\begin{aligned} \dot{\mathbf{v}}_l &= \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L) \\ &= -\sum_{j=1}^m [\mathbf{S} \mathbf{X}^T \mathbf{r}]_j (\mathbf{v}_1^T [\mathbf{U}_1]_{\cdot, j} \otimes \dots \otimes \mathbf{v}_{l-1}^T [\mathbf{U}_{l-1}]_{\cdot, j} \otimes [\mathbf{U}_l]_{\cdot, j} \otimes \mathbf{v}_{l+1}^T [\mathbf{U}_{l+1}]_{\cdot, j} \otimes \dots \otimes \mathbf{v}_L^T [\mathbf{U}_L]_{\cdot, j}) \\ &= -\sum_{j=1}^m [\mathbf{S} \mathbf{X}^T \mathbf{r}]_j \left(\prod_{k \neq l} [\mathbf{U}_k^T \mathbf{v}_k]_j \right) [\mathbf{U}_l]_{\cdot, j}. \end{aligned}$$

Left-multiplying \mathbf{U}_l^H (the conjugate transpose of \mathbf{U}_l) to both sides, we get

$$\mathbf{U}_l^H \dot{\mathbf{v}}_l = -\mathbf{S} \mathbf{X}^T \mathbf{r} \odot \prod_{k \neq l}^{\odot} \mathbf{U}_k^T \mathbf{v}_k, \quad (20)$$

where \prod^{\odot} denotes the product using entry-wise multiplication \odot .

Now consider the rate of growth for the absolute value squared of the j -th component of $\mathbf{U}_l^T \mathbf{v}_l$:

$$\begin{aligned} \frac{d}{dt} |[\mathbf{U}_l^T \mathbf{v}_l]_j|^2 &= \frac{d}{dt} [\mathbf{U}_l^T \mathbf{v}_l]_j [\mathbf{U}_l^T \mathbf{v}_l]_j^* = \frac{d}{dt} [\mathbf{U}_l^T \mathbf{v}_l]_j [\mathbf{U}_l^H \mathbf{v}_l]_j \\ &= [\mathbf{U}_l^T \dot{\mathbf{v}}_l]_j [\mathbf{U}_l^H \mathbf{v}_l]_j + [\mathbf{U}_l^H \dot{\mathbf{v}}_l]_j [\mathbf{U}_l^T \mathbf{v}_l]_j \\ &= 2 \operatorname{Re} ([\mathbf{U}_l^H \dot{\mathbf{v}}_l]_j [\mathbf{U}_l^T \mathbf{v}_l]_j) \\ &= 2 \operatorname{Re} \left(-[\mathbf{S} \mathbf{X}^T \mathbf{r}]_j \prod_{k=1}^L [\mathbf{U}_k^T \mathbf{v}_k]_j \right) \\ &= \frac{d}{dt} |[\mathbf{U}_{l'}^T \mathbf{v}_{l'}]_j|^2 \quad \text{for any } l' \in [L], \end{aligned}$$

so for any $j \in [m]$, the squared absolute value of the j -th components in $\mathbf{U}_l^T \mathbf{v}_l$ grow at the same rate for each layer $l \in [L]$. This means that the gap between any two different layers stays constant for all $t \geq 0$. Combining this with our conditions on initial directions, we have

$$\begin{aligned} |[\mathbf{U}_l^T \mathbf{v}_l(t)]_j|^2 - |[\mathbf{U}_L^T \mathbf{v}_L(t)]_j|^2 &= |[\mathbf{U}_l^T \mathbf{v}_l(0)]_j|^2 - |[\mathbf{U}_L^T \mathbf{v}_L(0)]_j|^2 \\ &= \alpha^2 |[\mathbf{U}_l^T \bar{\mathbf{v}}_l]_j|^2 - \alpha^2 |[\mathbf{U}_L^T \bar{\mathbf{v}}_L]_j|^2 \geq \alpha^2 \lambda, \end{aligned} \quad (21)$$

for any $j \in [m]$, $l \in [L-1]$, and $t \geq 0$. This inequality also implies

$$|[\mathbf{U}_l^T \mathbf{v}_l(t)]_j|^2 \geq |[\mathbf{U}_L^T \mathbf{v}_L(t)]_j|^2 + \alpha^2 \lambda \geq \alpha^2 \lambda. \quad (22)$$

Let us now consider the time derivative of $\mathcal{L}(\Theta(t))$. We have the following chain of upper bounds on the time derivative:

$$\begin{aligned}
 \frac{d}{dt}\mathcal{L}(\Theta(t)) &= \nabla_{\Theta}\mathcal{L}(\Theta(t))^T\dot{\Theta}(t) = -\|\nabla_{\Theta}\mathcal{L}(\Theta(t))\|_2^2 \\
 &\leq -\|\nabla_{\mathbf{v}_L}\mathcal{L}(\Theta(t))\|_2^2 = -\|\dot{\mathbf{v}}_L(t)\|_2^2 \\
 &\stackrel{(a)}{\leq} -\|\mathbf{U}_L^H\dot{\mathbf{v}}_L(t)\|_2^2 \stackrel{(b)}{=} -\left\|\mathbf{S}\mathbf{X}^T\mathbf{r}(t) \odot \prod_{k \neq L}^{\odot} \mathbf{U}_k^T \mathbf{v}_k(t)\right\|_2^2 \\
 &= -\sum_{j=1}^m |[\mathbf{S}\mathbf{X}^T\mathbf{r}(t)]_j|^2 \prod_{k \neq L} |[U_k^T \mathbf{v}_k(t)]_j|^2 \\
 &\stackrel{(c)}{\leq} -\alpha^{2L-2}\lambda^{L-1} \sum_{j=1}^m |[\mathbf{S}\mathbf{X}^T\mathbf{r}(t)]_j|^2 \\
 &= -\alpha^{2L-2}\lambda^{L-1} \|\mathbf{S}\mathbf{X}^T\mathbf{r}(t)\|_2^2 \\
 &\stackrel{(d)}{\leq} -\alpha^{2L-2}\lambda^{L-1} s_{\min}(\mathbf{S})^2 \|\mathbf{X}^T\mathbf{r}(t)\|_2^2, \tag{23}
 \end{aligned}$$

where (a) used the fact that $\|\dot{\mathbf{v}}_L(t)\|_2^2 \geq \|\mathbf{U}_L\mathbf{U}_L^H\dot{\mathbf{v}}_L(t)\|_2^2$ because it is a projection onto a subspace, and $\|\mathbf{U}_L\mathbf{U}_L^H\dot{\mathbf{v}}_L(t)\|_2^2 = \|\mathbf{U}_L^H\dot{\mathbf{v}}_L(t)\|_2^2$ because $\mathbf{U}_L^H\mathbf{U}_L = \mathbf{I}_{k_L}$; (b) is due to (20); (c) is due to (22); and (d) used the fact that $\mathbf{S} \in \mathbb{C}^{m \times d}$ is a matrix that has full column rank, so for any $\mathbf{z} \in \mathbb{C}^d$, we can use $\|\mathbf{S}\mathbf{z}\|_2 \geq s_{\min}(\mathbf{S})\|\mathbf{z}\|_2$ where $s_{\min}(\mathbf{S})$ is the minimum singular value of \mathbf{S} .

We now prove a lower bound on the quantity $\|\mathbf{X}^T\mathbf{r}(t)\|_2^2$. Recall from Section 2.1 the definition of $[\mathbf{r}(t)]_i = -y_i \exp(-y_i f(\mathbf{x}_i; \Theta(t)))$ for classification problems. Also, recall the assumption that the dataset is linearly separable, which means that there exists a unit vector $\mathbf{z} \in \mathbb{R}^d$ such that

$$y_i \mathbf{x}_i^T \mathbf{z} \geq \gamma > 0$$

holds for all $i \in [n]$, for some $\gamma > 0$. Using these,

$$\begin{aligned}
 \|\mathbf{X}^T\mathbf{r}(t)\|_2^2 &= \left\| \sum_{i=1}^n y_i \mathbf{x}_i \exp(-y_i f(\mathbf{x}_i; \Theta(t))) \right\|_2^2 \\
 &\geq \left[\mathbf{z}^T \sum_{i=1}^n y_i \mathbf{x}_i \exp(-y_i f(\mathbf{x}_i; \Theta(t))) \right]^2 \\
 &\geq \gamma^2 \left[\sum_{i=1}^n \exp(-y_i f(\mathbf{x}_i; \Theta(t))) \right]^2 = \gamma^2 \mathcal{L}(\Theta(t))^2.
 \end{aligned}$$

Combining this with (23), we get

$$\frac{d}{dt}\mathcal{L}(\Theta(t)) \leq -\alpha^{2L-2}\lambda^{L-1} s_{\min}(\mathbf{S})^2 \gamma^2 \mathcal{L}(\Theta(t))^2,$$

which implies

$$\mathcal{L}(\Theta(t)) \leq \frac{\mathcal{L}(\Theta(0))}{1 + \alpha^{2L-2}\lambda^{L-1} s_{\min}(\mathbf{S})^2 \gamma^2 t}.$$

Therefore, $\mathcal{L}(\Theta(t)) \rightarrow 0$ as $t \rightarrow \infty$.

H.1.2. CHARACTERIZING THE LIMIT DIRECTION

Since we proved that $\mathcal{L}(\Theta(t)) \rightarrow 0$, the argument in the proof of Theorem 1 applies to this case, and shows that the parameters \mathbf{v}_l converge in direction and align with $\dot{\mathbf{v}}_l = -\nabla_{\mathbf{v}_l}\mathcal{L}(\Theta)$. Let

$\mathbf{v}_l^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{v}_l(t)}{\|\mathbf{v}_l(t)\|_2}$ be the limit direction of \mathbf{v}_l .

The remaining steps of the proof are as follows. We first prove that $\mathbf{S}\mathbf{X}^T\mathbf{r}(t)$ converges in direction \mathbf{u}^∞ . Using this \mathbf{u}^∞ , we derive a number of conditions that has to be satisfied by the limit directions of the parameters. Finally, we compare these conditions with the KKT conditions of the minimization problem, and finish the proof.

By Assumption 1, we have

$$\begin{aligned} f(\mathbf{x}; \Theta) &= \mathbf{M}(\mathbf{x}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_L) = \sum_{j=1}^m [\mathbf{S}\mathbf{x}]_j \prod_{l=1}^L [\mathbf{U}_l^T \mathbf{v}_l]_j \\ &= \left[\sum_{j=1}^m \left(\prod_{l=1}^L [\mathbf{U}_l^T \mathbf{v}_l]_j \right) [\mathbf{S}]_{j\cdot} \right] \mathbf{x} = \mathbf{x}^T \mathbf{S}^T \left(\prod_{l \in [L]}^\odot \mathbf{U}_l^T \mathbf{v}_l \right) = \mathbf{x}^T \mathbf{S}^T \boldsymbol{\rho}. \end{aligned}$$

Here, we defined $\boldsymbol{\rho} := \prod_{l \in [L]}^\odot \mathbf{U}_l^T \mathbf{v}_l \in \mathbb{C}^m$. Since the linear coefficients must be real, we have $\mathbf{S}^T \boldsymbol{\rho} \in \mathbb{R}^d$ for any real \mathbf{v}_l 's. Since \mathbf{v}_l 's converge in direction, $\boldsymbol{\rho}$ also converges in direction, to $\boldsymbol{\rho}^\infty := \prod_{l \in [L]}^\odot \mathbf{U}_l^T \mathbf{v}_l^\infty$. So we can express the limit direction of $\beta(\Theta)$ as

$$\beta(\Theta^\infty) \propto \mathbf{S}^T \left(\prod_{l \in [L]}^\odot \mathbf{U}_l^T \mathbf{v}_l^\infty \right) = \mathbf{S}^T \boldsymbol{\rho}^\infty. \quad (24)$$

From (20) and alignment of \mathbf{v}_l and $\dot{\mathbf{v}}_l$, we have

$$\lim_{t \rightarrow \infty} \mathbf{U}_l^H \mathbf{v}_l(t) = \lim_{t \rightarrow \infty} (\mathbf{U}_l^T \mathbf{v}_l(t))^* \propto - \lim_{t \rightarrow \infty} \mathbf{S}\mathbf{X}^T \mathbf{r}(t) \odot \prod_{k \neq l}^\odot \mathbf{U}_k^T \mathbf{v}_k(t). \quad (25)$$

Since all vectors $\mathbf{U}_l^T \mathbf{v}_l(t)$ converge in direction, the term $\mathbf{S}\mathbf{X}^T \mathbf{r}(t)$ should also converge in direction. Let $\mathbf{u}^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{S}\mathbf{X}^T \mathbf{r}(t)}{\|\mathbf{S}\mathbf{X}^T \mathbf{r}(t)\|_2}$. One can use the same argument as in Appendix G.2, more specifically (16) and (17), to show that \mathbf{u}^∞ can be written as

$$\mathbf{u}^\infty \propto \lim_{t \rightarrow \infty} \mathbf{S} \sum_{i=1}^n \mathbf{x}_i [\mathbf{r}(t)]_i = -\mathbf{S} \sum_{i=1}^n \nu_i y_i \mathbf{x}_i, \quad (26)$$

where $\nu_i \geq 0$ for all $i \in [n]$, and $\nu_j = 0$ for \mathbf{x}_j 's that are *not* support vectors, i.e., those satisfying $y_j \mathbf{x}_j^T \mathbf{S}^T \boldsymbol{\rho}^\infty > \min_{i \in [n]} y_i \mathbf{x}_i^T \mathbf{S}^T \boldsymbol{\rho}^\infty$.

Using \mathbf{u}^∞ , we can rewrite (25) as

$$\mathbf{U}_l^H \mathbf{v}_l^\infty \propto -\mathbf{u}^\infty \odot \prod_{k \neq l}^\odot \mathbf{U}_k^T \mathbf{v}_k^\infty,$$

for all $l \in [L]$. Element-wise multiplying $\mathbf{U}_l^T \mathbf{v}_l^\infty$ to both sides gives

$$\mathbf{U}_l^T \mathbf{v}_l^\infty \odot \mathbf{U}_l^H \mathbf{v}_l^\infty = |\mathbf{U}_l^T \mathbf{v}_l^\infty|^{\odot 2} \propto -\mathbf{u}^\infty \odot \prod_{k \in [L]}^\odot \mathbf{U}_k^T \mathbf{v}_k^\infty = -\mathbf{u}^\infty \odot \boldsymbol{\rho}^\infty, \quad (27)$$

where $\mathbf{a}^{\odot b}$ denotes element-wise b -th power of the vector \mathbf{a} . Since the LHS of (27) is a positive real number, we have

$$\arg(|[\mathbf{U}_l^T \mathbf{v}_l]_j|^2) = 0 = \arg([-u^\infty]_j) + \arg([\boldsymbol{\rho}^\infty]_j), \quad (28)$$

so using this, (27) becomes

$$|\mathbf{U}_l^T \mathbf{v}_l^\infty|^{\odot 2} \propto |\mathbf{u}^\infty| \odot |\boldsymbol{\rho}^\infty|. \quad (29)$$

Now element-wise multiply (29) for all $l \in [L]$, then we get

$$|\boldsymbol{\rho}^\infty|^{\odot 2} \propto |\mathbf{u}^\infty|^{\odot L} \odot |\boldsymbol{\rho}^\infty|^{\odot L}. \quad (30)$$

A close look at (30) reveals that if $L \geq 2$, $\boldsymbol{\rho}^\infty$ and \mathbf{u}^∞ must satisfy that

$$|[\boldsymbol{\rho}^\infty]_j| \neq 0 \implies |[\mathbf{u}^\infty]_j| \propto |[\boldsymbol{\rho}^\infty]_j|^{\frac{2}{L}-1}, \quad (31)$$

for all $j \in [m]$. There is another condition that has to be satisfied when $L = 2$:

$$|[\boldsymbol{\rho}^\infty]_j| = 0, |[\boldsymbol{\rho}^\infty]_{j'}| \neq 0 \implies |[\mathbf{u}^\infty]_j| \leq |[\mathbf{u}^\infty]_{j'}|, \quad (32)$$

for any $j, j' \in [m]$; let us prove why. First, consider the time derivative of $[\boldsymbol{\rho}]_j = [\mathbf{U}_1^T \mathbf{v}_1]_j [\mathbf{U}_2^T \mathbf{v}_2]_j$.

$$\begin{aligned} \frac{d}{dt} [\boldsymbol{\rho}(t)]_j &= [\mathbf{U}_1^T \mathbf{v}_1(t)]_j \frac{d}{dt} [\mathbf{U}_2^T \mathbf{v}_2(t)]_j + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j \frac{d}{dt} [\mathbf{U}_1^T \mathbf{v}_1(t)]_j \\ &\stackrel{(a)}{=} -[\mathbf{S}\mathbf{X}^T \mathbf{r}(t)]_j^* (|[\mathbf{U}_1^T \mathbf{v}_1(t)]_j|^2 + |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2), \end{aligned} \quad (33)$$

where (a) used (20). Now consider

$$\frac{\left| \frac{d}{dt} [\boldsymbol{\rho}(t)]_j \right|}{\|\mathbf{S}\mathbf{X}^T \mathbf{r}(t)\|_2 |[\boldsymbol{\rho}(t)]_j|} = \frac{|\mathbf{S}\mathbf{X}^T \mathbf{r}(t)|_j}{\|\mathbf{S}\mathbf{X}^T \mathbf{r}(t)\|_2} \frac{|[\mathbf{U}_1^T \mathbf{v}_1(t)]_j|^2 + |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2}{|[\boldsymbol{\rho}(t)]_j|}. \quad (34)$$

We want to compare this quantity for different $j, j' \in [m]$. Before we do that, we take a look at the last term in the RHS of (34). Recall from (21) that

$$|[\mathbf{U}_1^T \mathbf{v}_1(t)]_j|^2 = |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2 + |[\mathbf{U}_1^T \mathbf{v}_1(0)]_j|^2 - |[\mathbf{U}_2^T \mathbf{v}_2(0)]_j|^2. \quad (35)$$

For simplicity, let $\delta_j := |[\mathbf{U}_1^T \mathbf{v}_1(0)]_j|^2 - |[\mathbf{U}_2^T \mathbf{v}_2(0)]_j|^2$. Then, we can use (35) and $|[\boldsymbol{\rho}(t)]_j| = |[\mathbf{U}_1^T \mathbf{v}_1(t)]_j| |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|$ to show that

$$\begin{aligned} \frac{|[\mathbf{U}_1^T \mathbf{v}_1(t)]_j|^2 + |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2}{|[\boldsymbol{\rho}(t)]_j|} &= \frac{2|[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2 + \delta_j}{|[\mathbf{U}_2^T \mathbf{v}_2(t)]_j| \sqrt{|[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2 + \delta_j}} \geq 2, \\ \lim_{t \rightarrow \infty} \frac{|[\mathbf{U}_1^T \mathbf{v}_1(t)]_j|^2 + |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2}{|[\boldsymbol{\rho}(t)]_j|} &= 2 \quad \text{if} \quad \lim_{t \rightarrow \infty} |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j| = \infty. \end{aligned}$$

Suppose that there exists $j \in [m]$ that satisfies $|[\boldsymbol{\rho}^\infty]_j| = 0$ but $|[\mathbf{u}^\infty]_j| > |[\mathbf{u}^\infty]_{j'}|$, for some $j' \in [m]$ satisfying $|[\boldsymbol{\rho}^\infty]_{j'}| \neq 0$. Note that having $|[\boldsymbol{\rho}^\infty]_j| = 0$ and $|[\boldsymbol{\rho}^\infty]_{j'}| \neq 0$ implies that $|[\boldsymbol{\rho}(t)]_{j'}| \rightarrow \infty$ and $\frac{|[\boldsymbol{\rho}(t)]_j|}{|[\boldsymbol{\rho}(t)]_{j'}|} \rightarrow 0$. We now want to compare the ratio of (34) for j and j' . First, note that

$$\lim_{t \rightarrow \infty} \frac{|[\mathbf{S}\mathbf{X}^T \mathbf{r}(t)]_j| / \|\mathbf{S}\mathbf{X}^T \mathbf{r}(t)\|_2}{|[\mathbf{S}\mathbf{X}^T \mathbf{r}(t)]_{j'}| / \|\mathbf{S}\mathbf{X}^T \mathbf{r}(t)\|_2} = \frac{|[\mathbf{u}^\infty]_j|}{|[\mathbf{u}^\infty]_{j'}|} > 1. \quad (36)$$

Next, using $\frac{|[\boldsymbol{\rho}(t)]_j|}{|[\boldsymbol{\rho}(t)]_{j'}|} \rightarrow 0$ and the fact that $x \mapsto \frac{2x^2 + \delta}{x\sqrt{x^2 + \delta}}$ is a decreasing function of $x \geq 0$ for any $\delta > 0$, we have

$$\frac{(|[\mathbf{U}_1^T \mathbf{v}_1(t)]_j|^2 + |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|^2) / |[\boldsymbol{\rho}(t)]_j|}{(|[\mathbf{U}_1^T \mathbf{v}_1(t)]_{j'}|^2 + |[\mathbf{U}_2^T \mathbf{v}_2(t)]_{j'}|^2) / |[\boldsymbol{\rho}(t)]_{j'}|} \geq 1, \quad (37)$$

for any $t \geq t_0$, when t_0 is large enough. Combining (36) and (37) to compare the ratio of (34) for j and j' , we get that there exists some $t'_0 \geq 0$ such that for any $t \geq t'_0$, we have

$$\frac{\left| \frac{d}{dt} [\boldsymbol{\rho}(t)]_j \right| / |[\boldsymbol{\rho}(t)]_j|}{\left| \frac{d}{dt} [\boldsymbol{\rho}(t)]_{j'} \right| / |[\boldsymbol{\rho}(t)]_{j'}|} > 1. \quad (38)$$

This implies that the ratio of the absolute value of time derivative of $[\boldsymbol{\rho}(t)]_j$ to the absolute value of current value of $[\boldsymbol{\rho}(t)]_j$ is strictly bigger than that of $[\boldsymbol{\rho}(t)]_{j'}$. Moreover, we saw in (33) that the phase of $\frac{d}{dt} [\boldsymbol{\rho}(t)]_j$ converges to that of $-\mathbf{u}^\infty_j$. Since this holds for all $t \geq t'_0$, (38) results in a growth of $|[\boldsymbol{\rho}(t)]_j|$ that is exponentially faster than that of $|[\boldsymbol{\rho}(t)]_{j'}|$, so $[\boldsymbol{\rho}(t)]_j$ becomes a dominant component in $\boldsymbol{\rho}(t)$ as $t \rightarrow \infty$. This contradicts that $[\boldsymbol{\rho}^\infty]_j = 0$, hence the condition (32) has to be satisfied.

So far, we have characterized a number of conditions (26), (28), (31), (32) that have to be satisfied by the limit directions \mathbf{u}^∞ and $\boldsymbol{\rho}^\infty$ of $\mathbf{X}^T \mathbf{r}$ and $\boldsymbol{\rho}$. We now consider the following optimization problem and prove that these conditions are in fact the KKT conditions of the optimization problem. Consider

$$\underset{\boldsymbol{\rho} \in \mathbb{C}^m}{\text{minimize}} \quad \|\boldsymbol{\rho}\|_{2/L} \quad \text{subject to} \quad y_i \mathbf{x}_i^T \mathbf{S}^T \boldsymbol{\rho} \geq 1, \quad \forall i \in [n]. \quad (39)$$

The KKT conditions of this problem are

$$\partial \|\boldsymbol{\rho}\|_{2/L} \ni \mathbf{S}^* \sum_{i=1}^n \mu_i y_i \mathbf{x}_i, \quad \text{and} \quad \mu_i \geq 0, \quad \mu_i (1 - y_i \mathbf{x}_i^T \mathbf{S}^T \boldsymbol{\rho}) = 0 \quad \text{for all } i \in [n],$$

where μ_1, \dots, μ_n are the dual variables. The symbol $\partial \|\cdot\|_{2/L}$ denotes the (local) subdifferential of the $\ell_{2/L}$ norm, which can be written as

$$\partial \|\boldsymbol{\rho}\|_1 = \{ \mathbf{u} \in \mathbb{C}^m \mid |[\mathbf{u}]_j| \leq 1 \text{ for all } j \in [m], \text{ and } [\boldsymbol{\rho}]_j \neq 0 \implies [\mathbf{u}]_j = \exp(\sqrt{-1} \arg([\boldsymbol{\rho}]_j)) \},$$

if $L = 2$ (in this case $\partial \|\boldsymbol{\rho}\|_1$ is the global subdifferential), and

$$\partial \|\boldsymbol{\rho}\|_{2/L} = \left\{ \mathbf{u} \in \mathbb{C}^m \mid [\boldsymbol{\rho}]_j \neq 0 \implies [\mathbf{u}]_j = \frac{2}{L} |[\boldsymbol{\rho}]_j|^{\frac{2}{L}-1} \exp(\sqrt{-1} \arg([\boldsymbol{\rho}]_j)) \right\},$$

if $L > 2$. By replacing μ_i 's with ν_i 's defined in (26), we can check from (26), (28), (31), (32) that the that $\boldsymbol{\rho}^\infty$ and \mathbf{u}^∞ satisfy the KKT conditions up to scaling. Therefore, by (24), $\boldsymbol{\beta}(\boldsymbol{\Theta}(t))$ converges in direction aligned with $\mathbf{S}^T \boldsymbol{\rho}^\infty$, where $\boldsymbol{\rho}^\infty$ is again aligned with a stationary point (global minimum in case of $L = 2$) of the optimization problem (39).

If \mathbf{S} is invertible, we can get $\mathbf{S}^{-T} \boldsymbol{\beta}(\boldsymbol{\Theta}^\infty) \propto \boldsymbol{\rho}^\infty$. Plugging this into the optimization problem (39) gives the last statement of the theorem.

H.2. Proof of Corollary 9

It suffices to prove that linear diagonal networks satisfy Assumption 1, with $\mathbf{S} = \mathbf{I}_d$. The proof is very straightforward, since $\mathbf{M}_{\text{diag}}(\mathbf{x}) \in \mathbb{R}^{d \times \dots \times d}$ has $[\mathbf{M}_{\text{diag}}(\mathbf{x})]_{j,j,\dots,j} = [\mathbf{x}]_j$ while all the remaining entries are zero. It is straightforward to verify that $\mathbf{M}_{\text{diag}}(\mathbf{x})$ satisfies Assumption 1 with $\mathbf{S} = \mathbf{U}_1 = \dots = \mathbf{U}_L = \mathbf{I}_d$. A direct substitution into Theorem 2 gives the corollary.

H.3. Proof of Corollary 10

For full-length convolutional networks ($k_1 = \dots = k_L = d$), we will prove that they satisfy Assumption 1 with $\mathbf{S} = d^{\frac{L-1}{2}} \mathbf{F}$ and $\mathbf{U}_1 = \dots = \mathbf{U}_L = \mathbf{F}^*$, where $\mathbf{F} \in \mathbb{C}^{d \times d}$ is the matrix of discrete Fourier transform basis $[\mathbf{F}]_{j,k} = \frac{1}{\sqrt{d}} \exp(-\frac{\sqrt{-1} \cdot 2\pi(j-1)(k-1)}{d})$ and \mathbf{F}^* is the complex conjugate of \mathbf{F} .

For simplicity of notation, define $\psi = \exp(-\frac{\sqrt{-1} \cdot 2\pi}{d})$. With such matrices \mathbf{S} and $\mathbf{U}_1, \dots, \mathbf{U}_L$, we can write $\mathbf{M}(\mathbf{x})$ as

$$\begin{aligned} \mathbf{M}(\mathbf{x}) &= \sum_{j=1}^d [\mathbf{S}\mathbf{x}]_j ([\mathbf{U}_1]_{\cdot,j} \otimes [\mathbf{U}_2]_{\cdot,j} \otimes \dots \otimes [\mathbf{U}_L]_{\cdot,j}) \\ &= \sum_{j=1}^d \left[d^{\frac{L-2}{2}} \sum_{k=1}^d [\mathbf{x}]_k \psi^{(j-1)(k-1)} \right] \begin{bmatrix} \psi^0 / \sqrt{d} \\ \psi^{-(j-1)} / \sqrt{d} \\ \psi^{-2(j-1)} / \sqrt{d} \\ \vdots \\ \psi^{-(d-1)(j-1)} / \sqrt{d} \end{bmatrix}^{\otimes L}, \end{aligned}$$

where $\mathbf{a}^{\otimes L}$ denotes the L -times tensor product of \mathbf{a} . We will show that $\mathbf{M}(\mathbf{x}) = \mathbf{M}_{\text{conv}}(\mathbf{x})$.

For any $j_1, \dots, j_L \in [d]$,

$$\begin{aligned} [\mathbf{M}(\mathbf{x})]_{j_1, \dots, j_L} &= \frac{1}{d} \sum_{l=1}^d \left[\sum_{k=1}^d [\mathbf{x}]_k \psi^{(l-1)(k-1)} \right] \psi^{-(l-1)(\sum_{q=1}^L j_q - L)} \\ &= \frac{1}{d} \sum_{k=1}^d [\mathbf{x}]_k \sum_{l=1}^d \psi^{(l-1)(k-1 - \sum_{q=1}^L j_q + L)}. \end{aligned}$$

Recall that

$$\sum_{l=1}^d \psi^{(l-1)(k-1 - \sum_{q=1}^L j_q + L)} = \begin{cases} d & \text{if } k-1 - \sum_{q=1}^L j_q + L \text{ is a multiple of } d, \\ 0 & \text{otherwise.} \end{cases}$$

Using this, we have

$$\begin{aligned} [\mathbf{M}(\mathbf{x})]_{j_1, \dots, j_L} &= \frac{1}{d} \sum_{k=1}^d [\mathbf{x}]_k \sum_{l=1}^d \psi^{(l-1)(k-1 - \sum_{q=1}^L j_q + L)} \\ &= [\mathbf{x}]_{\sum_{q=1}^L j_q - L + 1 \bmod d} = [\mathbf{M}_{\text{conv}}(\mathbf{x})]_{j_1, \dots, j_L}. \end{aligned}$$

Hence, linear full-length convolutional networks satisfy Assumption 1 with $\mathbf{S} = d^{\frac{L-1}{2}} \mathbf{F}$. A direct substitution into Theorem 2 and then using the fact that $|\mathbf{F}\mathbf{z}]_j| = |[\mathbf{F}^*\mathbf{z}]_j|$ for any real vector $\mathbf{z} \in \mathbb{R}^d$ gives the corollary.

Appendix I. Proofs of Theorem 4, Corollaries 11, 12 & 13, and Lemma 3

I.1. Proof of Lemma 3

In this subsection, we restate Lemma 3 and prove it.

Lemma 15 Consider the system of ODEs, where $p, q : \mathbb{R} \rightarrow \mathbb{R}$:

$$\dot{p} = p^{L-2}q, \quad \dot{q} = p^{L-1}, \quad p(0) = 1, \quad q(0) = 0.$$

Then, the solutions $p_L(t)$ and $q_L(t)$ are continuous on their maximal interval of existence of the form $(-c, c) \subset \mathbb{R}$ for some $c \in (0, \infty]$. Define $h_L(t) = p_L(t)^{L-1}q_L(t)$; then, $h_L(t)$ is odd and strictly increasing, satisfying $\lim_{t \uparrow c} h_L(t) = \infty$ and $\lim_{t \downarrow -c} h_L(t) = -\infty$.

Proof First, continuity of $p(t)$ and $q(t)$ is straightforward because they are the solution of the ODE. Next, define $\tilde{p}(t) = p(-t)$ and $\tilde{q}(t) = -q(-t)$. Then, one can show that \tilde{p} and \tilde{q} are also the solution of the ODE because

$$\begin{aligned} \frac{d}{dt}\tilde{p}(t) &= \frac{d}{dt}p(-t) = -\dot{p}(-t) = -p(-t)^{L-2}q(-t) = \tilde{p}(t)^{L-2}\tilde{q}(t), \\ \frac{d}{dt}\tilde{q}(t) &= -\frac{d}{dt}q(-t) = \dot{q}(-t) = p(-t)^{L-1} = \tilde{p}(t)^{L-1}. \end{aligned}$$

However, by the Picard-Lindelöf theorem, the solution has to be unique; this means that $p(t) = \tilde{p}(t) = p(-t)$ and $q(t) = \tilde{q}(t) = -q(-t)$, which proves that p is even and q is odd and also implies that the domain of p and q has to be of the form $(-c, c)$ (i.e. symmetric around the origin) and $h = p^{L-1}q$ is odd.

To show that h is strictly increasing, it suffices to show that p and q are both strictly increasing on $[0, c)$. To this end, we show that $p(t) \geq 1$ for all $t \in [0, c)$. First, due to the initial condition $p(0) = 1$ and continuity of p , there exists $\epsilon_1 > 0$ such that $p(t) > 0$ for all $t \in [0, \epsilon_1) =: I_1$. This implies that $\dot{q}(t) = p(t)^{L-1} > 0$ for $t \in I_1$, so q is strictly increasing on I_1 . Since $q(0) = 0$, we have $q(t) > 0$ for $t \in I_1$, which then implies that $\dot{p}(t) = p(t)^{L-2}q(t) > 0$. Therefore, p is also strictly increasing on I_1 ; this then means $p(t) \geq 1$ for $t \in [0, \epsilon_1]$ because $p(0) = 1$. Now, due to $p(\epsilon_1) \geq 1$ and continuity of p , there exists $\epsilon_2 > \epsilon_1$ such that $p(t) > 0$ for all $t \in [\epsilon_1, \epsilon_2) =: I_2$. Using the argument above for I_2 results in $p(t) \geq 1$ for $t \in [0, \epsilon_2]$. Repeating this until the end of the domain, we can show that $p(t) \geq 1$ holds for all $t \in [0, c)$. By $p \geq 1$, we have $\dot{q} = p^{L-1} \geq 1$ on $[0, c)$, so q is strictly increasing on $[0, c)$. Also, $q(t) > 0$ on $(0, c)$, so $\dot{p} = p^{L-2}q > 0$ on $(0, c)$ and p is also strictly increasing on $[0, c)$. This proves that h is strictly increasing on $[0, c)$, and also on $(-c, c)$ by oddity of h .

Finally, it is left to show $\lim_{t \uparrow c} h(t) = \infty$ and $\lim_{t \downarrow -c} h(t) = -\infty$. If $c < \infty$, then this together with monotonicity implies that the limits hold. To see why, suppose $c < \infty$ and $\lim_{t \uparrow c} h(t) < \infty$. Then, p and q can be extended beyond $t \geq c$, which contradicts the fact that $(-c, c)$ is the maximal interval of existence of the solution. Next, consider the case $c = \infty$. From $p(t) \geq 1$, we have $\dot{q}(t) \geq 1$ for $t \geq 0$. This implies that $q(t) \geq t$ for $t \geq 0$. Now, $\dot{p}(t) \geq p(t)^{L-2}q(t) \geq t$, which gives $p(t) \geq \frac{t^2}{2} + 1$ for $t \geq 0$. Therefore, we have

$$\lim_{t \rightarrow \infty} h(t) = \lim_{t \rightarrow \infty} p(t)^{L-1}q(t) \geq \lim_{t \rightarrow \infty} \left(\frac{t^2}{2} + 1 \right)^{L-1} t = \infty,$$

hence finishing the proof. ■

I.2. Proof of Theorem 4

I.2.1. CONVERGENCE OF LOSS TO ZERO

We first show that given the conditions on initialization, the training loss $\mathcal{L}(\Theta(t))$ converges to zero. Recall from Section 2.1 that

$$\dot{\mathbf{v}}_l = -\nabla_{\mathbf{v}_l} \mathcal{L}(\Theta) = \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L).$$

Applying the structure (5) in Assumption 1, we get

$$\begin{aligned} \dot{\mathbf{v}}_l &= \mathbf{M}(-\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \dots, \mathbf{v}_{l-1}, \mathbf{I}_{k_l}, \mathbf{v}_{l+1}, \dots, \mathbf{v}_L) \\ &= -\sum_{j=1}^m [\mathbf{S} \mathbf{X}^T \mathbf{r}]_j (\mathbf{v}_1^T [\mathbf{U}_1]_{\cdot, j} \otimes \dots \otimes \mathbf{v}_{l-1}^T [\mathbf{U}_{l-1}]_{\cdot, j} \otimes [\mathbf{U}_l]_{\cdot, j} \otimes \mathbf{v}_{l+1}^T [\mathbf{U}_{l+1}]_{\cdot, j} \otimes \dots \otimes \mathbf{v}_L^T [\mathbf{U}_L]_{\cdot, j}) \\ &= -\sum_{j=1}^m [\mathbf{S} \mathbf{X}^T \mathbf{r}]_j \left(\prod_{k \neq l} [\mathbf{U}_k^T \mathbf{v}_k]_j \right) [\mathbf{U}_l]_{\cdot, j}. \end{aligned}$$

Left-multiplying \mathbf{U}_l^T to both sides, we get

$$\mathbf{U}_l^T \dot{\mathbf{v}}_l = -\mathbf{S} \mathbf{X}^T \mathbf{r} \odot \prod_{k \neq l}^{\odot} \mathbf{U}_k^T \mathbf{v}_k, \quad (40)$$

where \prod^{\odot} denotes the product using entry-wise multiplication \odot .

Now consider the rate of growth for the second power of the j -th component of $\mathbf{U}_l^T \mathbf{v}_l$:

$$\frac{d}{dt} [\mathbf{U}_l^T \mathbf{v}_l]_j^2 = 2[\mathbf{U}_l^T \dot{\mathbf{v}}_l]_j [\mathbf{U}_l^T \mathbf{v}_l]_j = -2[\mathbf{S} \mathbf{X}^T \mathbf{r}]_j \prod_{k=1}^L [\mathbf{U}_k^T \mathbf{v}_k]_j = \frac{d}{dt} |[\mathbf{U}_{l'}^T \mathbf{v}_{l'}]_j|^2$$

for any $l' \in [L]$. Thus, for any $j \in [m]$, the second power of the j -th components in $\mathbf{U}_l^T \mathbf{v}_l$ grow at the same rate for each layer $l \in [L]$. This means that the gap between any two different layers stays constant for all $t \geq 0$. Combining this with our conditions on initial directions, we have

$$[\mathbf{U}_l^T \mathbf{v}_l(t)]_j^2 - [\mathbf{U}_L^T \mathbf{v}_L(t)]_j^2 = [\mathbf{U}_l^T \mathbf{v}_l(0)]_j^2 - [\mathbf{U}_L^T \mathbf{v}_L(0)]_j^2 = \alpha^2 [\bar{\eta}]_j^2 \geq \alpha^2 \lambda,$$

for any $j \in [m]$, $l \in [L-1]$, and $t \geq 0$. This inequality also implies

$$[\mathbf{U}_l^T \mathbf{v}_l(t)]_j^2 \geq [\mathbf{U}_L^T \mathbf{v}_L(t)]_j^2 + \alpha^2 \lambda \geq \alpha^2 \lambda. \quad (41)$$

Let us now consider the time derivative of $\mathcal{L}(\Theta(t))$. We have the following chain of upper bounds on the time derivative:

$$\begin{aligned} \frac{d}{dt} \mathcal{L}(\Theta(t)) &= \nabla_{\Theta} \mathcal{L}(\Theta(t))^T \dot{\Theta}(t) = -\|\nabla_{\Theta} \mathcal{L}(\Theta(t))\|_2^2 \\ &\leq -\|\nabla_{\mathbf{v}_L} \mathcal{L}(\Theta(t))\|_2^2 = -\|\dot{\mathbf{v}}_L(t)\|_2^2 \\ &\stackrel{(a)}{\leq} -\|\mathbf{U}_L^T \dot{\mathbf{v}}_L(t)\|_2^2 \stackrel{(b)}{=} -\left\| \mathbf{S} \mathbf{X}^T \mathbf{r}(t) \odot \prod_{k \neq L}^{\odot} \mathbf{U}_k^T \mathbf{v}_k(t) \right\|_2^2 \\ &= -\sum_{j=1}^m [\mathbf{S} \mathbf{X}^T \mathbf{r}(t)]_j^2 \prod_{k \neq L} [\mathbf{U}_k^T \mathbf{v}_k(t)]_j^2 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(c)}{\leq} -\alpha^{2L-2}\lambda^{L-1}\sum_{j=1}^m[\mathbf{S}\mathbf{X}^T\mathbf{r}(t)]_j^2 \\
 &= -\alpha^{2L-2}\lambda^{L-1}\|\mathbf{S}\mathbf{X}^T\mathbf{r}(t)\|_2^2 \\
 &\stackrel{(d)}{\leq} -\alpha^{2L-2}\lambda^{L-1}s_{\min}(\mathbf{S})^2s_{\min}(\mathbf{X})^2\|\mathbf{r}(t)\|_2^2, \\
 &= -2\alpha^{2L-2}\lambda^{L-1}s_{\min}(\mathbf{S})^2s_{\min}(\mathbf{X})^2\mathcal{L}(\Theta(t)),
 \end{aligned} \tag{42}$$

where (a) used the fact that $\|\dot{\mathbf{v}}_L(t)\|_2^2 \geq \|\mathbf{U}_L\mathbf{U}_L^T\dot{\mathbf{v}}_L(t)\|_2^2$ because it is a projection onto a subspace, and $\|\mathbf{U}_L\mathbf{U}_L^T\dot{\mathbf{v}}_L(t)\|_2^2 = \|\mathbf{U}_L^T\dot{\mathbf{v}}_L(t)\|_2^2$ because $\mathbf{U}_L^T\mathbf{U}_L = \mathbf{I}_{k_L}$; (b) is due to (40); (c) is due to (41); and (d) used the fact that $\mathbf{S} \in \mathbb{R}^{m \times d}$ and $\mathbf{X}^T \in \mathbb{R}^{d \times n}$ are matrices that have full column rank, so for any $\mathbf{z} \in \mathbb{C}^n$, we can use $\|\mathbf{S}\mathbf{X}^T\mathbf{z}\|_2 \geq s_{\min}(\mathbf{S})s_{\min}(\mathbf{X})\|\mathbf{z}\|_2$ where $s_{\min}(\cdot)$ denotes the minimum singular value of a matrix.

From (42), we get

$$\mathcal{L}(\Theta(t)) \leq \mathcal{L}(\Theta(0)) \exp(-2\alpha^{2L-2}\lambda^{L-1}s_{\min}(\mathbf{S})^2s_{\min}(\mathbf{X})^2t), \tag{43}$$

so that $\mathcal{L}(\Theta(t)) \rightarrow 0$ as $t \rightarrow \infty$.

I.2.2. CHARACTERIZING THE LIMIT POINT

Now, we move on to characterize the limit points of the gradient flow. First, by defining a ‘‘transformed’’ version of the parameters $\boldsymbol{\eta}_l(t) := \mathbf{U}_l^T\mathbf{v}_l(t)$ and using (40), one can define an equivalent system of ODEs:

$$\begin{aligned}
 \dot{\boldsymbol{\eta}}_l &= -\mathbf{S}\mathbf{X}^T\mathbf{r} \odot \prod_{k \neq l}^{\odot} \boldsymbol{\eta}_k \text{ for } l \in [L], \\
 \boldsymbol{\eta}_l(0) &= \alpha\bar{\boldsymbol{\eta}} \text{ for } l \in [L-1], \quad \boldsymbol{\eta}_L(0) = \mathbf{0}.
 \end{aligned} \tag{44}$$

Using Lemma 3, it is straightforward to verify that the solution to (44) has the following form. For odd L , we have

$$\begin{aligned}
 \boldsymbol{\eta}_l(t) &= \alpha\bar{\boldsymbol{\eta}} \odot p_L \left(-\alpha^{L-2}|\bar{\boldsymbol{\eta}}|^{\odot L-2} \odot \mathbf{S}\mathbf{X}^T \int_0^t \mathbf{r}(\tau) d\tau \right) \text{ for } l \in [L-1], \\
 \boldsymbol{\eta}_L(t) &= \alpha|\bar{\boldsymbol{\eta}}| \odot q_L \left(-\alpha^{L-2}|\bar{\boldsymbol{\eta}}|^{\odot L-2} \odot \mathbf{S}\mathbf{X}^T \int_0^t \mathbf{r}(\tau) d\tau \right).
 \end{aligned} \tag{45}$$

Similarly, for even L , the solution for (44) satisfies

$$\begin{aligned}
 \boldsymbol{\eta}_l(t) &= \alpha\bar{\boldsymbol{\eta}} \odot p_L \left(-\alpha^{L-2}\bar{\boldsymbol{\eta}}^{\odot L-2} \odot \mathbf{S}\mathbf{X}^T \int_0^t \mathbf{r}(\tau) d\tau \right) \text{ for } l \in [L-1], \\
 \boldsymbol{\eta}_L(t) &= \alpha\bar{\boldsymbol{\eta}} \odot q_L \left(-\alpha^{L-2}\bar{\boldsymbol{\eta}}^{\odot L-2} \odot \mathbf{S}\mathbf{X}^T \int_0^t \mathbf{r}(\tau) d\tau \right).
 \end{aligned} \tag{46}$$

Now that we know how the solutions $\boldsymbol{\eta}_l$ look like, let us see how these relate to the linear coefficients of the network. By Assumption 1, we have

$$f(\mathbf{x}; \Theta) = \mathbf{M}(\mathbf{x}) \odot (\mathbf{v}_1, \dots, \mathbf{v}_L) = \sum_{j=1}^m [\mathbf{S}\mathbf{x}]_j \prod_{l=1}^L [\mathbf{U}_l^T\mathbf{v}_l]_j$$

$$= \left[\sum_{j=1}^m \left(\prod_{l=1}^L [\boldsymbol{\eta}]_j \right) [\mathbf{S}]_{j\cdot} \right] \mathbf{x} = \mathbf{x}^T \mathbf{S}^T \left(\prod_{l \in [L]} \boldsymbol{\eta}_l \right) = \mathbf{x}^T \mathbf{S}^T \boldsymbol{\rho}.$$

Here, we defined $\boldsymbol{\rho} := \prod_{l \in [L]} \boldsymbol{\eta}_l \in \mathbb{R}^m$. Therefore, the linear coefficients of the network can be written as $\boldsymbol{\beta}(\boldsymbol{\Theta}(t)) = \mathbf{S}^T \boldsymbol{\rho}(t)$. From the solutions (45) and (46), we can write

$$\boldsymbol{\rho}(t) = \prod_{i=1}^L \boldsymbol{\eta}_i(t) = \alpha^L |\bar{\boldsymbol{\eta}}|^{\odot L} \odot h_L \left(-\alpha^{L-2} |\bar{\boldsymbol{\eta}}|^{\odot L-2} \odot \mathbf{S} \mathbf{X}^T \int_0^t \mathbf{r}(\tau) d\tau \right),$$

where $h_L := p_L^{L-1} q_L$, defined in Lemma 3. By the convergence of the loss to zero (43), we have $\lim_{t \rightarrow \infty} \mathbf{X} \boldsymbol{\beta}(\boldsymbol{\Theta}(t)) = \mathbf{y}$. Therefore,

$$\underbrace{\mathbf{X} \mathbf{S}^T \left(\alpha^L |\bar{\boldsymbol{\eta}}|^{\odot L} \odot h_L \left(-\alpha^{L-2} |\bar{\boldsymbol{\eta}}|^{\odot L-2} \odot \mathbf{S} \mathbf{X}^T \int_0^\infty \mathbf{r}(\tau) d\tau \right) \right)}_{=: \boldsymbol{\rho}^\infty} = \mathbf{y}. \quad (47)$$

Next, we will show that $\boldsymbol{\rho}^\infty$ is in fact the solution of the following optimization problem

$$\underset{\boldsymbol{\rho} \in \mathbb{R}^m}{\text{minimize}} \quad Q_{L,\alpha,\bar{\boldsymbol{\eta}}}(\boldsymbol{\rho}) \quad \text{subject to} \quad \mathbf{X} \mathbf{S}^T \boldsymbol{\rho} = \mathbf{y}, \quad (48)$$

where $Q_{L,\alpha,\bar{\boldsymbol{\eta}}} : \mathbb{R}^m \rightarrow \mathbb{R}$ is a norm-like function defined using $H_L(t) := \int_0^t h_L^{-1}(\tau) d\tau$:

$$Q_{L,\alpha,\bar{\boldsymbol{\eta}}}(\boldsymbol{\rho}) = \alpha^2 \sum_{j=1}^m [\bar{\boldsymbol{\eta}}]_j^2 H_L \left(\frac{[\boldsymbol{\rho}]_j}{\alpha^L |[\bar{\boldsymbol{\eta}}]_j|^L} \right).$$

Note that the KKT conditions for (48) are

$$\mathbf{X} \mathbf{S}^T \boldsymbol{\rho} = \mathbf{y}, \quad \nabla_{\boldsymbol{\rho}} Q_{L,\alpha,\bar{\boldsymbol{\eta}}}(\boldsymbol{\rho}) = \mathbf{S} \mathbf{X}^T \boldsymbol{\nu},$$

for some $\boldsymbol{\nu} \in \mathbb{R}^n$. It is clear from (47) that $\boldsymbol{\rho}^\infty$ satisfies the first condition (primal feasibility), so let us check the other one. Through a straightforward calculation, we get

$$\nabla_{\boldsymbol{\rho}} Q_{L,\alpha,\bar{\boldsymbol{\eta}}}(\boldsymbol{\rho}) = \alpha^{2-L} |\bar{\boldsymbol{\eta}}|^{\odot 2-L} \odot h_L^{-1} \left(\alpha^{-L} |\bar{\boldsymbol{\eta}}|^{\odot (-L)} \odot \boldsymbol{\rho} \right).$$

Equating this with $\mathbf{S} \mathbf{X}^T \boldsymbol{\nu}$ gives

$$\begin{aligned} & \alpha^{2-L} |\bar{\boldsymbol{\eta}}|^{\odot 2-L} \odot h_L^{-1} \left(\alpha^{-L} |\bar{\boldsymbol{\eta}}|^{\odot (-L)} \odot \boldsymbol{\rho} \right) = \mathbf{S} \mathbf{X}^T \boldsymbol{\nu} \\ \Leftrightarrow & h_L^{-1} \left(\alpha^{-L} |\bar{\boldsymbol{\eta}}|^{\odot (-L)} \odot \boldsymbol{\rho} \right) = \alpha^{L-2} |\bar{\boldsymbol{\eta}}|^{\odot L-2} \odot \mathbf{S} \mathbf{X}^T \boldsymbol{\nu} \\ \Leftrightarrow & \boldsymbol{\rho} = \alpha^L |\bar{\boldsymbol{\eta}}|^{\odot L} \odot h_L \left(\alpha^{L-2} |\bar{\boldsymbol{\eta}}|^{\odot L-2} \odot \mathbf{S} \mathbf{X}^T \boldsymbol{\nu} \right). \end{aligned}$$

Hence, by setting $\boldsymbol{\nu} = -\int_0^\infty \mathbf{r}(\tau) d\tau$, $\boldsymbol{\rho}^\infty$ satisfies this condition as well. Also, if \mathbf{S} is invertible, we can substitute $\boldsymbol{\rho} = \mathbf{S}^{-T} \mathbf{z}$ to (48) to get the last statement of the theorem. This finishes the proof.

I.3. Proof of Corollary 11

The proof is a direct consequence of the fact that Assumption 1 holds with $\mathbf{S} = \mathbf{U}_1 = \dots = \mathbf{U}_L = \mathbf{I}_d$ for linear diagonal networks. Hence, the proof is the same as Corollary 9, proved in Appendix H.2.

I.4. Proof of Corollary 12

We start by showing the DFT of a real and even vector is also real and even. Suppose that $\mathbf{x} \in \mathbb{R}^d$ is real and even. First,

$$\begin{aligned} [\mathbf{F}\mathbf{x}]_j &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \exp\left(-\frac{\sqrt{-1} \cdot 2\pi(j-1)(k-1)}{d}\right) \\ &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \cos\left(-\frac{2\pi(j-1)(k-1)}{d}\right) + \frac{\sqrt{-1}}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \sin\left(-\frac{2\pi(j-1)(k-1)}{d}\right) \\ &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \cos\left(-\frac{2\pi(j-1)(k-1)}{d}\right) \in \mathbb{R}, \end{aligned}$$

for all $j \in [d]$. To prove that $\mathbf{F}\mathbf{x}$ is even, for $j = 0, \dots, \lfloor \frac{d-3}{2} \rfloor$, we have

$$\begin{aligned} [\mathbf{F}\mathbf{x}]_{j+2} &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \cos\left(-\frac{2\pi(j+1)(k-1)}{d}\right) \\ &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \cos\left(2\pi(k-1) - \frac{2\pi(j+1)(k-1)}{d}\right) \\ &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \cos\left(\frac{2\pi(d-j-1)(k-1)}{d}\right) \\ &= \frac{1}{\sqrt{d}} \sum_{k=1}^d [\mathbf{x}]_k \cos\left(-\frac{2\pi(d-j-1)(k-1)}{d}\right) \\ &= [\mathbf{F}\mathbf{x}]_{d-j}. \end{aligned}$$

It is proved in Appendix H.3 that linear full-length convolutional networks ($k_1 = \dots = k_L = d$) satisfy Assumption 1 with $\mathbf{S} = d^{\frac{L-1}{2}} \mathbf{F}$ and $\mathbf{U}_1 = \dots = \mathbf{U}_L = \mathbf{F}^*$, where $\mathbf{F} \in \mathbb{C}^{d \times d}$ is the matrix of discrete Fourier transform basis $[\mathbf{F}]_{j,k} = \frac{1}{\sqrt{d}} \exp(-\frac{\sqrt{-1} \cdot 2\pi(j-1)(k-1)}{d})$ and \mathbf{F}^* is the complex conjugate of \mathbf{F} .

The proof of convergence of loss to zero in Appendix I.2.1 is written for real matrices $\mathbf{S}, \mathbf{U}_1, \dots, \mathbf{U}_L$, but we can actually apply the same argument as in Appendix H.1.1 and prove that the loss converges to zero, even in the case where $\mathbf{S}, \mathbf{U}_1, \dots, \mathbf{U}_L$ are complex.

Next, since \mathbf{U}_l 's are complex, we can write the system of ODE as (see (20) for its derivation)

$$\mathbf{F}\dot{\mathbf{w}}_l = -d^{\frac{L-1}{2}} \mathbf{F}\mathbf{X}^T \mathbf{r} \odot \prod_{k \neq l}^{\odot} \mathbf{F}^* \mathbf{w}_k, \quad (49)$$

Since all data points \mathbf{x}_i and initialization $\mathbf{w}_l(0)$ are real and even, we have that $\mathbf{F}\mathbf{X}^T\mathbf{r}$ is real and even, and $\mathbf{F}^*\mathbf{w}_l(0) = \mathbf{F}\mathbf{w}_l(0)$'s are real and even. By (49), we see that the time derivatives of $\mathbf{F}\mathbf{w}_l$ are also real and even. Thus, the parameters $\mathbf{w}_l(t)$ are all real and even for all $t \geq 0$. From this observation, we can define $\boldsymbol{\eta}_l(t) := \mathbf{F}\mathbf{w}_l(t)$, $\bar{\boldsymbol{\eta}} := \mathbf{F}\bar{\mathbf{w}}$, and $\mathbf{S} := d^{\frac{L-1}{2}} \text{Re}(\mathbf{F})$, which are all real by the even symmetry. Then, starting from (44), the proof goes through.

I.5. Proof of Corollary 13

Since the sensor matrices $\mathbf{A}_1, \dots, \mathbf{A}_n$ commute, they are simultaneously diagonalizable with a real unitary matrix $\mathbf{U} \in \mathbb{R}^{d \times d}$, i.e., $\mathbf{U}^T \mathbf{A}_i \mathbf{U}$'s are diagonal matrices. From the deep matrix sensing problem (13), we can compute $\nabla_{\mathbf{W}_l} \partial \mathcal{L}$, which gives the gradient flow dynamics of \mathbf{W}_l .

$$\dot{\mathbf{W}}_l = -\nabla_{\mathbf{W}_l} \partial \mathcal{L} = -\mathbf{W}_{l-1}^T \cdots \mathbf{W}_1^T \left(\sum_{i=1}^n r_i \mathbf{A}_i \right) \mathbf{W}_L^T \cdots \mathbf{W}_{l+1}^T,$$

where $r_i = \langle \mathbf{A}_i, \mathbf{W}_1 \cdots \mathbf{W}_L \rangle - y_i$ is the residual for the i -th sensor matrix. If all we left-multiply \mathbf{U}^T and right-multiply \mathbf{U} to both sides, we get

$$\mathbf{U}^T \dot{\mathbf{W}}_l \mathbf{U} = -\mathbf{U}^T \mathbf{W}_{l-1}^T \mathbf{U} \cdots \mathbf{U}^T \mathbf{W}_1^T \mathbf{U} \left(\sum_{i=1}^n r_i \mathbf{U}^T \mathbf{A}_i \mathbf{U} \right) \mathbf{U}^T \mathbf{W}_L^T \mathbf{U} \cdots \mathbf{U}^T \mathbf{W}_{l+1}^T \mathbf{U}. \quad (50)$$

If $\mathbf{U}^T \mathbf{W}_k^T \mathbf{U}$ is a diagonal matrix for all $k \neq l$, then $\mathbf{U}^T \dot{\mathbf{W}}_l \mathbf{U}$ is also a diagonal matrix. Note also that, since $\mathbf{W}_l(0) = \alpha \mathbf{I}_d = \alpha \mathbf{U} \mathbf{U}^T$ for $l \in [L-1]$, the product $\mathbf{U}^T \mathbf{W}_l \mathbf{U}$ is a diagonal matrix at initialization. These observations imply that $\mathbf{W}_l(t)$'s are all diagonalizable with \mathbf{U} for all $t \geq 0$.

Now, define $\mathbf{v}_l(t) = \text{eig}(\mathbf{W}_l(t))$, i.e., $\mathbf{U}^T \mathbf{W}_l \mathbf{U} = \text{diag}(\mathbf{v}_l)$. Also, let $\mathbf{x}_i = \text{eig}(\mathbf{A}_i)$. Then, (50) can be written as

$$\dot{\mathbf{v}}_l = -\left(\sum_{i=1}^n r_i \mathbf{x}_i \right) \odot \prod_{k \neq l}^{\odot} \mathbf{v}_k.$$

Therefore, this is equivalent to the regression problem with linear diagonal networks, initialized at $\mathbf{v}_l(0) = \alpha \mathbf{1}$ for $l \in [L-1]$ and $\mathbf{v}_L(0) = \mathbf{0}$. Given this equivalence, Corollary 13 can be implied from Corollary 11.

Appendix J. Proofs of Theorem 5 and Corollary 14

J.1. Proof of Theorem 5

J.1.1. CONVERGENCE OF LOSS TO ZERO

Since Theorem 5 does not assume the existence of $t_0 \geq 0$ satisfying $\mathcal{L}(\boldsymbol{\Theta}(t_0)) < 1$, we need to first show that given the conditions on initialization, the training loss $\mathcal{L}(\boldsymbol{\Theta}(t))$ converges to zero. Since $L = 2$ and $\mathbf{M}(\mathbf{x}) = \mathbf{U}_1 \text{diag}(\mathbf{s}) \mathbf{U}_2^T$, we can write the gradient flow dynamics from Section 2.1 as

$$\begin{aligned} \dot{\mathbf{v}}_1 &= -\mathbf{M}(\mathbf{X}^T \mathbf{r}) \circ (\mathbf{I}_{k_1}, \mathbf{v}_2) = -r \mathbf{U}_1 \text{diag}(\mathbf{s}) \mathbf{U}_2^T \mathbf{v}_2, \\ \dot{\mathbf{v}}_2 &= -\mathbf{M}(\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \mathbf{I}_{k_2}) = -r \mathbf{U}_2 \text{diag}(\mathbf{s}) \mathbf{U}_1^T \mathbf{v}_1, \end{aligned} \quad (51)$$

where $r(t) = -y \exp(-yf(\mathbf{x}; \boldsymbol{\Theta}(t)))$ is the residual of the data point (\mathbf{x}, y) . From (51) we get

$$\mathbf{U}_1^T \dot{\mathbf{v}}_1 = -r \mathbf{s} \odot \mathbf{U}_2^T \mathbf{v}_2, \quad \mathbf{U}_2^T \dot{\mathbf{v}}_2 = -r \mathbf{s} \odot \mathbf{U}_1^T \mathbf{v}_1. \quad (52)$$

Now consider the rate of growth for the j -th component of $\mathbf{U}_1^T \mathbf{v}_1$ squared:

$$\frac{d}{dt} [\mathbf{U}_1^T \mathbf{v}_1]_j^2 = 2[\mathbf{U}_1^T \mathbf{v}_1]_j [\mathbf{U}_1^T \dot{\mathbf{v}}_1]_j = -2r[\mathbf{s}]_j [\mathbf{U}_1^T \mathbf{v}_1]_j [\mathbf{U}_2^T \mathbf{v}_2]_j = \frac{d}{dt} [\mathbf{U}_2^T \mathbf{v}_2]_j^2. \quad (53)$$

So for any $j \in [m]$, $[\mathbf{U}_1^T \mathbf{v}_1]_j^2$ and $[\mathbf{U}_2^T \mathbf{v}_2]_j^2$ grow at the same rate. This means that the gap between the two layers stays constant for all $t \geq 0$. Combining this with our conditions on initial directions,

$$\begin{aligned} [\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 - [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 &= [\mathbf{U}_1^T \mathbf{v}_1(0)]_j^2 - [\mathbf{U}_2^T \mathbf{v}_2(0)]_j^2 \\ &= \alpha^2 [\mathbf{U}_1^T \bar{\mathbf{v}}_1]_j^2 - \alpha^2 [\mathbf{U}_2^T \bar{\mathbf{v}}_2]_j^2 \geq \alpha^2 \lambda, \end{aligned} \quad (54)$$

for any $j \in [m]$ and $t \geq 0$. This inequality implies

$$[\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 \geq [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 + \alpha^2 \lambda \geq \alpha^2 \lambda. \quad (55)$$

Let us now consider the time derivative of $\mathcal{L}(\Theta(t))$. We have the following chain of upper bounds on the time derivative:

$$\begin{aligned} \frac{d}{dt} \mathcal{L}(\Theta(t)) &= \nabla_{\Theta} \mathcal{L}(\Theta(t))^T \dot{\Theta}(t) = -\|\nabla_{\Theta} \mathcal{L}(\Theta(t))\|_2^2 \\ &\leq -\|\nabla_{\mathbf{v}_2} \mathcal{L}(\Theta(t))\|_2^2 = -\|\dot{\mathbf{v}}_2(t)\|_2^2 \\ &\stackrel{(a)}{\leq} -\|\mathbf{U}_2^T \dot{\mathbf{v}}_2(t)\|_2^2 \stackrel{(b)}{=} -r(t)^2 \|\mathbf{s} \odot \mathbf{U}_1^T \mathbf{v}_1(t)\|_2^2 \\ &= -r(t)^2 \sum_{j=1}^m [\mathbf{s}]_j^2 [\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 \\ &\stackrel{(c)}{\leq} -\alpha^2 \lambda r(t)^2 \sum_{j=1}^m [\mathbf{s}]_j^2 \\ &= -\alpha^2 \lambda \|\mathbf{s}\|_2^2 \mathcal{L}(\Theta(t))^2, \end{aligned}$$

where (a) used the fact that $\|\dot{\mathbf{v}}_2(t)\|_2^2 \geq \|\mathbf{U}_2 \mathbf{U}_2^T \dot{\mathbf{v}}_2(t)\|_2^2$ because it is a projection onto a subspace, and $\|\mathbf{U}_2 \mathbf{U}_2^T \dot{\mathbf{v}}_2(t)\|_2^2 = \|\mathbf{U}_2^T \dot{\mathbf{v}}_2(t)\|_2^2$ because $\mathbf{U}_2^T \mathbf{U}_2 = \mathbf{I}_{k_2}$; (b) is due to (52); (c) is due to (55). From this, we get

$$\mathcal{L}(\Theta(t)) \leq \frac{\mathcal{L}(\Theta(0))}{1 + \alpha^2 \lambda \|\mathbf{s}\|_2^2 t}.$$

Therefore, $\mathcal{L}(\Theta(t)) \rightarrow 0$ as $t \rightarrow \infty$.

J.1.2. CHARACTERIZING THE LIMIT DIRECTION

Since we proved that $\mathcal{L}(\Theta(t)) \rightarrow 0$, the argument in the proof of Theorem 1 applies to this case, and shows that the parameters \mathbf{v}_l converge in direction and align with $\dot{\mathbf{v}}_l = -\nabla_{\mathbf{v}_l} \mathcal{L}(\Theta)$. Let $\mathbf{v}_l^\infty := \lim_{t \rightarrow \infty} \frac{\mathbf{v}_l(t)}{\|\mathbf{v}_l(t)\|_2}$ be the limit direction of \mathbf{v}_l . As done in the proof of Theorem 2, define $\boldsymbol{\rho}(t) = \mathbf{U}_1^T \mathbf{v}_1(t) \odot \mathbf{U}_2^T \mathbf{v}_2(t)$ and $\boldsymbol{\rho}^\infty = \mathbf{U}_1^T \mathbf{v}_1^\infty \odot \mathbf{U}_2^T \mathbf{v}_2^\infty$.

It follows from $r(t) = -y \exp(-yf(\mathbf{x}; \Theta(t)))$ that we have $\text{sign}(r(t)) = -\text{sign}(y)$. Using this, (52), and alignment of \mathbf{v}_l and $\dot{\mathbf{v}}_l$, we have

$$\mathbf{U}_1^T \mathbf{v}_1^\infty \propto y \mathbf{s} \odot \mathbf{U}_2^T \mathbf{v}_2^\infty, \quad \mathbf{U}_2^T \mathbf{v}_2^\infty \propto y \mathbf{s} \odot \mathbf{U}_1^T \mathbf{v}_1^\infty. \quad (56)$$

Element-wise multiplying $\mathbf{U}_l^T \mathbf{v}_l^\infty$ to both sides gives

$$(\mathbf{U}_1^T \mathbf{v}_1^\infty)^{\odot 2} \propto \mathbf{y} \mathbf{s} \odot \boldsymbol{\rho}^\infty, \quad (\mathbf{U}_2^T \mathbf{v}_2^\infty)^{\odot 2} \propto \mathbf{y} \mathbf{s} \odot \boldsymbol{\rho}^\infty. \quad (57)$$

Since the LHSs are positive and \mathbf{s} is positive, the following equations have to be satisfied for all $j \in [m]$:

$$\text{sign}(y) = \text{sign}([\boldsymbol{\rho}^\infty]_j). \quad (58)$$

Now, multiplying both sides of the two equations (57), we get

$$(\boldsymbol{\rho}^\infty)^{\odot 2} \propto \mathbf{s}^{\odot 2} \odot (\boldsymbol{\rho}^\infty)^{\odot 2}. \quad (59)$$

From (59), $\boldsymbol{\rho}^\infty$ must satisfy that

$$[\boldsymbol{\rho}^\infty]_j \neq 0, [\boldsymbol{\rho}^\infty]_{j'} \neq 0 \implies |[s]_j| = |[s]_{j'}|, \quad (60)$$

for all $j, j' \in [m]$. As in the proof of Theorem 2, there is another condition that has to be satisfied:

$$[\boldsymbol{\rho}^\infty]_j = 0, [\boldsymbol{\rho}^\infty]_{j'} \neq 0 \implies |[s]_j| \leq |[s]_{j'}|, \quad (61)$$

for any $j, j' \in [m]$; let us prove why. First, consider the time derivative of $[\boldsymbol{\rho}]_j = [\mathbf{U}_1^T \mathbf{v}_1]_j [\mathbf{U}_2^T \mathbf{v}_2]_j$.

$$\begin{aligned} \frac{d}{dt} [\boldsymbol{\rho}(t)]_j &= [\mathbf{U}_1^T \mathbf{v}_1(t)]_j \frac{d}{dt} [\mathbf{U}_2^T \mathbf{v}_2(t)]_j + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j \frac{d}{dt} [\mathbf{U}_1^T \mathbf{v}_1(t)]_j \\ &\stackrel{(a)}{=} -r(t) [s]_j ([\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2), \end{aligned}$$

where (a) used (52). Now consider

$$\frac{\left| \frac{d}{dt} [\boldsymbol{\rho}(t)]_j \right|}{|r(t)| |[\boldsymbol{\rho}(t)]_j|} = |[s]_j| \frac{[\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2}{|[\boldsymbol{\rho}(t)]_j|}. \quad (62)$$

We want to compare this quantity for different $j, j' \in [m]$. Before we do that, we take a look at the last term in the RHS of (62). Recall from (54) that

$$[\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 = [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 + [\mathbf{U}_1^T \mathbf{v}_1(0)]_j^2 - [\mathbf{U}_2^T \mathbf{v}_2(0)]_j^2. \quad (63)$$

For simplicity, let $\delta_j := [\mathbf{U}_1^T \mathbf{v}_1(0)]_j^2 - [\mathbf{U}_2^T \mathbf{v}_2(0)]_j^2$. Then, we can use (63) and $|[\boldsymbol{\rho}(t)]_j| = |[\mathbf{U}_1^T \mathbf{v}_1(t)]_j| |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j|$ to show that

$$\begin{aligned} \frac{[\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2}{|[\boldsymbol{\rho}(t)]_j|} &= \frac{2[\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 + \delta_j}{|[\mathbf{U}_2^T \mathbf{v}_2(t)]_j| \sqrt{[\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 + \delta_j}} \geq 2, \\ \lim_{t \rightarrow \infty} \frac{[\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2}{|[\boldsymbol{\rho}(t)]_j|} &= 2 \quad \text{if} \quad \lim_{t \rightarrow \infty} |[\mathbf{U}_2^T \mathbf{v}_2(t)]_j| = \infty. \end{aligned}$$

Suppose that there exists $j \in [m]$ that satisfies $[\boldsymbol{\rho}^\infty]_j = 0$ but $|[s]_j| > |[s]_{j'}|$, for some $j' \in [m]$ satisfying $[\boldsymbol{\rho}^\infty]_{j'} \neq 0$. Note that having $[\boldsymbol{\rho}^\infty]_j = 0$ and $[\boldsymbol{\rho}^\infty]_{j'} \neq 0$ implies that $|[\boldsymbol{\rho}(t)]_{j'}| \rightarrow \infty$ and

$\frac{|\rho(t)_j|}{|\rho(t)_{j'}|} \rightarrow 0$. We now want to compare the ratio of (62) for j and j' . Using $\frac{|\rho(t)_j|}{|\rho(t)_{j'}|} \rightarrow 0$ and the fact that $x \mapsto \frac{2x^2 + \delta}{x\sqrt{x^2 + \delta}}$ is a decreasing function of $x \geq 0$ for any $\delta > 0$, we have

$$\frac{([\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 + [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2) / |\rho(t)_j|}{([\mathbf{U}_1^T \mathbf{v}_1(t)]_{j'}^2 + [\mathbf{U}_2^T \mathbf{v}_2(t)]_{j'}^2) / |\rho(t)_{j'}|} \geq 1, \quad (64)$$

for any $t \geq t_0$, when t_0 is large enough. Combining $\frac{|\mathbf{s}_j|}{|\mathbf{s}_{j'}|} > 1$ and (64) to compare the ratio of (62) for j and j' , there exists some $t_0 \geq 0$ such that for any $t \geq t_0$, we have

$$\frac{\left| \frac{d}{dt} |\rho(t)_j| \right| / |\rho(t)_j|}{\left| \frac{d}{dt} |\rho(t)_{j'}| \right| / |\rho(t)_{j'}|} > 1. \quad (65)$$

This implies that the ratio of the absolute value of time derivative of $|\rho(t)_j|$ to the absolute value of current value of $|\rho(t)_j|$ is strictly bigger than that of $|\rho(t)_{j'}|$. Moreover, by the definition of $r(t)$, $\frac{d}{dt} |\rho(t)_j|$ does not change sign over time. Since this holds for all $t \geq t_0$, (65) results in a growth of $|\rho(t)_j|$ that is exponentially faster than that of $|\rho(t)_{j'}|$, so $|\rho(t)_j|$ becomes a dominant component in $\rho(t)$ as $t \rightarrow \infty$. This contradicts that $[\rho^\infty]_j = 0$, hence the condition (61) has to be satisfied.

So far, we have characterized some conditions (58), (60), (61) that have to be satisfied by the limit direction ρ^∞ of ρ . We now consider the following optimization problem and prove that these conditions are in fact the KKT conditions of the optimization problem. Consider

$$\underset{\rho \in \mathbb{R}^m}{\text{minimize}} \quad \|\rho\|_1 \quad \text{subject to} \quad \mathbf{y} \mathbf{s}^T \rho \geq 1. \quad (66)$$

The KKT condition of this problem is

$$\partial \|\rho\|_1 \ni \mathbf{y} \mathbf{s},$$

where the global subdifferential $\partial \|\cdot\|_1$ is defined as

$$\partial \|\rho\|_1 = \{ \mathbf{u} \in \mathbb{R}^m \mid |u_j| \leq 1 \text{ for all } j \in [m], \text{ and } [\rho]_j \neq 0 \implies [u]_j = \text{sign}([\rho]_j) \}.$$

We can check from (58), (60), (61) that the that ρ^∞ satisfies the KKT condition up to scaling.

Now, how do we characterize \mathbf{v}_1^∞ and \mathbf{v}_2^∞ in terms of ρ^∞ ? Let $\boldsymbol{\eta}_1^\infty := \mathbf{U}_1^T \mathbf{v}_1^\infty$ and $\boldsymbol{\eta}_2^\infty := \mathbf{U}_2^T \mathbf{v}_2^\infty$. Then, $\mathbf{v}_i^\infty = \mathbf{U}_i \boldsymbol{\eta}_i^\infty = \mathbf{U}_i \mathbf{U}_i^T \mathbf{v}_i^\infty$ holds because any component orthogonal to the column space of \mathbf{U}_i stays unchanged while the component in the column space of \mathbf{U}_i diverges to infinity. By (53), $|\boldsymbol{\eta}_1^\infty| = |\boldsymbol{\eta}_2^\infty| = |\rho^\infty|^{\odot 1/2}$. By (56), we have $\text{sign}(\boldsymbol{\eta}_1^\infty) = \text{sign}(\mathbf{y}) \odot \text{sign}(\boldsymbol{\eta}_2^\infty)$.

J.2. Proof of Corollary 14

The proof of Corollary 14 boils down to characterizing the SVD of $\mathbf{M}_{\text{conv}}(\mathbf{x})$.

J.2.1. THE $k_1 = 1$ CASE

First, it is straightforward to check that for $L = 2$ and $k_1 = 1$, we have

$$\boldsymbol{\beta}_{\text{conv}}(\boldsymbol{\Theta}_{\text{conv}}) = \mathbf{v}_1 \mathbf{v}_2.$$

For $k_1 = 1$, the data tensor is simply $\mathbf{M}_{\text{conv}}(\mathbf{x}) = \mathbf{x}^T$. Thus, we have $U_1 = 1$, $U_2 = \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$, and $\mathbf{s} = \|\mathbf{x}\|_2$. Substituting U_1 and U_2 to the theorem gives the condition on initial directions in Corollary 14. Also, the theorem implies us that the limit direction \mathbf{v}_2^∞ of \mathbf{v}_2 satisfies $\mathbf{v}_2^\infty \propto y\mathbf{v}_1^\infty \mathbf{x}$. Using this, it is easy to check that

$$\beta_{\text{conv}}(\Theta_{\text{conv}}^\infty) \propto \mathbf{v}_1^\infty \mathbf{v}_2^\infty \propto y\mathbf{x}.$$

J.2.2. THE $k_1 = 2$ CASE

First, it is straightforward to check that for $L = 2$ and $k_1 = 2$, we have

$$\beta_{\text{conv}}(\Theta_{\text{conv}}) = \begin{bmatrix} [\mathbf{v}_1]_1 & 0 & 0 & \cdots & 0 & [\mathbf{v}_1]_2 \\ [\mathbf{v}_1]_2 & [\mathbf{v}_1]_1 & 0 & \cdots & 0 & 0 \\ 0 & [\mathbf{v}_1]_2 & [\mathbf{v}_1]_1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & [\mathbf{v}_1]_1 & 0 \\ 0 & 0 & 0 & \cdots & [\mathbf{v}_1]_2 & [\mathbf{v}_1]_1 \end{bmatrix} \mathbf{v}_2. \quad (67)$$

For $k_1 = 2$, by definition, the data tensor is

$$\mathbf{M}_{\text{conv}}(\mathbf{x}) = \begin{bmatrix} \mathbf{x}^T \\ \overleftarrow{\mathbf{x}}^T \end{bmatrix},$$

and it is straightforward to check that the SVD of this matrix is

$$\mathbf{M}_{\text{conv}}(\mathbf{x}) = \begin{bmatrix} \mathbf{x}^T \\ \overleftarrow{\mathbf{x}}^T \end{bmatrix} = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix} \begin{bmatrix} \sqrt{\|\mathbf{x}\|_2^2 + \mathbf{x}^T \overleftarrow{\mathbf{x}}} & 0 \\ 0 & \sqrt{\|\mathbf{x}\|_2^2 - \mathbf{x}^T \overleftarrow{\mathbf{x}}} \end{bmatrix} \begin{bmatrix} \frac{\mathbf{x}^T + \overleftarrow{\mathbf{x}}^T}{\sqrt{2}\sqrt{\|\mathbf{x}\|_2^2 + \mathbf{x}^T \overleftarrow{\mathbf{x}}}} \\ \frac{\mathbf{x}^T - \overleftarrow{\mathbf{x}}^T}{\sqrt{2}\sqrt{\|\mathbf{x}\|_2^2 - \mathbf{x}^T \overleftarrow{\mathbf{x}}}} \end{bmatrix},$$

so

$$U_1 = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix}, U_2 = \begin{bmatrix} \frac{\mathbf{x} + \overleftarrow{\mathbf{x}}}{\sqrt{2}\sqrt{\|\mathbf{x}\|_2^2 + \mathbf{x}^T \overleftarrow{\mathbf{x}}}} & \frac{\mathbf{x} - \overleftarrow{\mathbf{x}}}{\sqrt{2}\sqrt{\|\mathbf{x}\|_2^2 - \mathbf{x}^T \overleftarrow{\mathbf{x}}}} \end{bmatrix}, \mathbf{s} = \begin{bmatrix} \sqrt{\|\mathbf{x}\|_2^2 + \mathbf{x}^T \overleftarrow{\mathbf{x}}} \\ \sqrt{\|\mathbf{x}\|_2^2 - \mathbf{x}^T \overleftarrow{\mathbf{x}}} \end{bmatrix}.$$

Substituting U_1 and U_2 to the theorem gives the conditions on initial directions. Also, note that the maximum singular value depends on the sign of $\mathbf{x}^T \overleftarrow{\mathbf{x}}$. Consider the optimization problem in the theorem statement:

$$\text{minimize}_{\boldsymbol{\rho} \in \mathbb{R}^m} \|\boldsymbol{\rho}\|_1 \quad \text{subject to} \quad y\mathbf{s}^T \boldsymbol{\rho} \geq 1.$$

If $\mathbf{x}^T \overleftarrow{\mathbf{x}} > 0$, then the solution $\boldsymbol{\rho}^\infty$ to this problem is in the direction of $[y \ 0]$. Therefore, the limit directions \mathbf{v}_1^∞ and \mathbf{v}_2^∞ will be of the form

$$\mathbf{v}_1^\infty \propto c_1 \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \mathbf{v}_2^\infty \propto c_2(\mathbf{x} + \overleftarrow{\mathbf{x}}),$$

where $\text{sign}(c_1) \text{sign}(c_2) = \text{sign}(y)$. Using (67), it is straightforward to check that

$$\beta_{\text{conv}}(\Theta_{\text{conv}}^\infty) \propto y \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 1 \\ 1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 1 & 1 \end{bmatrix} (\mathbf{x} + \overleftarrow{\mathbf{x}}) = y(2\mathbf{x} + \overleftarrow{\mathbf{x}} + \overrightarrow{\mathbf{x}}).$$

Similarly, if $\mathbf{x}^T \overleftarrow{\mathbf{x}} < 0$, then the solution ρ^∞ is in the direction of $[0 \ y]$. Using (67), we have

$$\beta_{\text{conv}}(\Theta_{\text{conv}}^\infty) \propto y \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & -1 \\ -1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & -1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & -1 & 1 \end{bmatrix} (\mathbf{x} - \overleftarrow{\mathbf{x}}) = y(2\mathbf{x} - \overleftarrow{\mathbf{x}} - \overrightarrow{\mathbf{x}}).$$

Appendix K. Proof of Theorem 6

K.1. Convergence of loss to zero

We first show that given the conditions on initialization, the training loss $\mathcal{L}(\Theta(t))$ converges to zero. Since $L = 2$ and $\mathbf{M}(\mathbf{x}) = \mathbf{U}_1 \text{diag}(\mathbf{s}) \mathbf{U}_2^T$, we can write the gradient flow dynamics from Section 2.1 as

$$\begin{aligned} \dot{\mathbf{v}}_1 &= -\mathbf{M}(\mathbf{X}^T \mathbf{r}) \circ (\mathbf{I}_{k_1}, \mathbf{v}_2) = -r \mathbf{U}_1 \text{diag}(\mathbf{s}) \mathbf{U}_2^T \mathbf{v}_2, \\ \dot{\mathbf{v}}_2 &= -\mathbf{M}(\mathbf{X}^T \mathbf{r}) \circ (\mathbf{v}_1, \mathbf{I}_{k_2}) = -r \mathbf{U}_2 \text{diag}(\mathbf{s}) \mathbf{U}_1^T \mathbf{v}_1, \end{aligned} \quad (68)$$

where $r(t) = f(\mathbf{x}; \Theta(t)) - y$ is the residual of the data point (\mathbf{x}, y) . From (68) we get

$$\mathbf{U}_1^T \dot{\mathbf{v}}_1 = -r \mathbf{s} \odot \mathbf{U}_2^T \mathbf{v}_2, \quad \mathbf{U}_2^T \dot{\mathbf{v}}_2 = -r \mathbf{s} \odot \mathbf{U}_1^T \mathbf{v}_1. \quad (69)$$

Now consider the rate of growth for the j -th component of $\mathbf{U}_1^T \mathbf{v}_1$ squared:

$$\frac{d}{dt} [\mathbf{U}_1^T \mathbf{v}_1]_j^2 = 2 [\mathbf{U}_1^T \mathbf{v}_1]_j [\mathbf{U}_1^T \dot{\mathbf{v}}_1]_j = -2r [\mathbf{s}]_j [\mathbf{U}_1^T \mathbf{v}_1]_j [\mathbf{U}_2^T \mathbf{v}_2]_j = \frac{d}{dt} [\mathbf{U}_2^T \mathbf{v}_2]_j^2.$$

So for any $j \in [m]$, $[\mathbf{U}_1^T \mathbf{v}_1]_j^2$ and $[\mathbf{U}_2^T \mathbf{v}_2]_j^2$ grow at the same rate. This means that the gap between the two layers stays constant for all $t \geq 0$. Combining this with our conditions on initial directions,

$$\begin{aligned} [\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 - [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 &= [\mathbf{U}_1^T \mathbf{v}_1(0)]_j^2 - [\mathbf{U}_2^T \mathbf{v}_2(0)]_j^2 \\ &= \alpha^2 [\mathbf{U}_1^T \bar{\mathbf{v}}_1]_j^2 - \alpha^2 [\mathbf{U}_2^T \bar{\mathbf{v}}_2]_j^2 \geq \alpha^2 \lambda, \end{aligned}$$

for any $j \in [m]$ and $t \geq 0$. This inequality implies

$$[\mathbf{U}_1^T \mathbf{v}_1(t)]_j^2 \geq [\mathbf{U}_2^T \mathbf{v}_2(t)]_j^2 + \alpha^2 \lambda \geq \alpha^2 \lambda. \quad (70)$$

Let us now consider the time derivative of $\mathcal{L}(\Theta(t))$. We have the following chain of upper bounds on the time derivative:

$$\begin{aligned}
\frac{d}{dt}\mathcal{L}(\Theta(t)) &= \nabla_{\Theta}\mathcal{L}(\Theta(t))^T\dot{\Theta}(t) = -\|\nabla_{\Theta}\mathcal{L}(\Theta(t))\|_2^2 \\
&\leq -\|\nabla_{\mathbf{v}_2}\mathcal{L}(\Theta(t))\|_2^2 = -\|\dot{\mathbf{v}}_2(t)\|_2^2 \\
&\stackrel{(a)}{\leq} -\|U_2^T\dot{\mathbf{v}}_2(t)\|_2^2 \stackrel{(b)}{=} -r(t)^2\|\mathbf{s}\odot U_1^T\mathbf{v}_1(t)\|_2^2 \\
&= -r(t)^2\sum_{j=1}^m[\mathbf{s}]_j^2[U_1^T\mathbf{v}_1(t)]_j^2 \\
&\stackrel{(c)}{\leq} -\alpha^2\lambda r(t)^2\sum_{j=1}^m[\mathbf{s}]_j^2 \\
&= -2\alpha^2\lambda\|\mathbf{s}\|_2^2\mathcal{L}(\Theta(t)),
\end{aligned}$$

where (a) used the fact that $\|\dot{\mathbf{v}}_2(t)\|_2^2 \geq \|U_2U_2^T\dot{\mathbf{v}}_2(t)\|_2^2$ because it is a projection onto a subspace, and $\|U_2U_2^T\dot{\mathbf{v}}_2(t)\|_2^2 = \|U_2^T\dot{\mathbf{v}}_2(t)\|_2^2$ because $U_2^TU_2 = \mathbf{I}_{k_2}$; (b) is due to (69); (c) is due to (70). From this, we get

$$\mathcal{L}(\Theta(t)) \leq \mathcal{L}(\Theta(0)) \exp(-2\alpha^2\lambda\|\mathbf{s}\|_2^2 t). \quad (71)$$

Therefore, $\mathcal{L}(\Theta(t)) \rightarrow 0$ as $t \rightarrow \infty$.

K.2. Characterizing the limit point

Now, we move on to characterize the limit points of the gradient flow. First, note that any changes made in \mathbf{v}_l over time are in the subspace spanned by the columns of U_l . Therefore, any component in the initialization $\mathbf{v}_l(0) = \alpha\bar{\mathbf{v}}_l$ that is orthogonal to the column space of U_l stays constant.

So, we can focus on the evolution of \mathbf{v}_l in the column space of U_l ; this can be done by defining a “transformed” version of the parameters $\boldsymbol{\eta}_l(t) := U_l^T\mathbf{v}_l(t)$ and using (69), one can define an equivalent system of ODEs:

$$\begin{aligned}
\dot{\boldsymbol{\eta}}_1 &= -r\mathbf{s}\odot\boldsymbol{\eta}_2, \quad \dot{\boldsymbol{\eta}}_2 = -r\mathbf{s}\odot\boldsymbol{\eta}_1, \\
\boldsymbol{\eta}_1(0) &= \alpha\bar{\boldsymbol{\eta}}_1, \quad \boldsymbol{\eta}_2(0) = \alpha\bar{\boldsymbol{\eta}}_2,
\end{aligned} \quad (72)$$

where $\bar{\boldsymbol{\eta}}_1 := U_1^T\bar{\mathbf{v}}_1$, $\bar{\boldsymbol{\eta}}_2 := U_2^T\bar{\mathbf{v}}_2$. It is straightforward to verify that the solution to (72) has the following form.

$$\begin{aligned}
\boldsymbol{\eta}_1(t) &= \alpha\bar{\boldsymbol{\eta}}_1 \odot \cosh\left(-\mathbf{s}\int_0^t r(\tau)d\tau\right) + \alpha\bar{\boldsymbol{\eta}}_2 \odot \sinh\left(-\mathbf{s}\int_0^t r(\tau)d\tau\right), \\
\boldsymbol{\eta}_2(t) &= \alpha\bar{\boldsymbol{\eta}}_1 \odot \sinh\left(-\mathbf{s}\int_0^t r(\tau)d\tau\right) + \alpha\bar{\boldsymbol{\eta}}_2 \odot \cosh\left(-\mathbf{s}\int_0^t r(\tau)d\tau\right).
\end{aligned} \quad (73)$$

By the convergence of the loss to zero (71), we have $\lim_{t \rightarrow \infty} f(\mathbf{x}; \Theta(t)) = y$. Note that $f(\mathbf{x}; \Theta(t))$ can be written as

$$\begin{aligned}
f(\mathbf{x}; \Theta(t)) &= \mathbf{M}(\mathbf{x}) \circ (\mathbf{v}_1(t), \mathbf{v}_2(t)) = \mathbf{v}_1(t)^T \mathbf{M}(\mathbf{x}) \mathbf{v}_2(t) \\
&= \mathbf{v}_1(t)^T U_1 \text{diag}(\mathbf{s}) U_2^T \mathbf{v}_2(t) = \mathbf{s}^T (\boldsymbol{\eta}_1(t) \odot \boldsymbol{\eta}_2(t)).
\end{aligned}$$

Therefore,

$$\begin{aligned}
 & \lim_{t \rightarrow \infty} f(\mathbf{x}; \Theta(t)) = \lim_{t \rightarrow \infty} \mathbf{s}^T (\boldsymbol{\eta}_1(t) \odot \boldsymbol{\eta}_2(t)) \\
 &= \alpha^2 \mathbf{s}^T \left[(\bar{\boldsymbol{\eta}}_1^{\odot 2} + \bar{\boldsymbol{\eta}}_2^{\odot 2}) \odot \cosh \left(-\mathbf{s} \int_0^\infty r(\tau) d\tau \right) \odot \sinh \left(-\mathbf{s} \int_0^\infty r(\tau) d\tau \right) \right. \\
 & \quad \left. + (\bar{\boldsymbol{\eta}}_1 \odot \bar{\boldsymbol{\eta}}_2) \odot \left(\cosh^{\odot 2} \left(-\mathbf{s} \int_0^\infty r(\tau) d\tau \right) + \sinh^{\odot 2} \left(-\mathbf{s} \int_0^\infty r(\tau) d\tau \right) \right) \right] \\
 &= \alpha^2 \mathbf{s}^T \left[\frac{\bar{\boldsymbol{\eta}}_1^{\odot 2} + \bar{\boldsymbol{\eta}}_2^{\odot 2}}{2} \odot \sinh \left(-2\mathbf{s} \int_0^\infty r(\tau) d\tau \right) + (\bar{\boldsymbol{\eta}}_1 \odot \bar{\boldsymbol{\eta}}_2) \odot \cosh \left(-2\mathbf{s} \int_0^\infty r(\tau) d\tau \right) \right] \\
 &= \alpha^2 \sum_{j=1}^m [\mathbf{s}]_j \left(\frac{[\bar{\boldsymbol{\eta}}_1]_j^2 + [\bar{\boldsymbol{\eta}}_2]_j^2}{2} \sinh(2[\mathbf{s}]_j \nu) + [\bar{\boldsymbol{\eta}}_1]_j [\bar{\boldsymbol{\eta}}_2]_j \cosh(2[\mathbf{s}]_j \nu) \right) \\
 &= y, \tag{74}
 \end{aligned}$$

where we defined $\nu := -\int_0^\infty r(\tau) d\tau$. Consider the function $\nu \mapsto a \sinh(\nu) + b \cosh(\nu)$. This is a strictly increasing function if $a > |b|$. Note also that

$$\frac{[\bar{\boldsymbol{\eta}}_1]_j^2 + [\bar{\boldsymbol{\eta}}_2]_j^2}{2} \geq |[\bar{\boldsymbol{\eta}}_1]_j [\bar{\boldsymbol{\eta}}_2]_j|, \tag{75}$$

which holds with equality if and only if $|[\bar{\boldsymbol{\eta}}_1]_j| = |[\bar{\boldsymbol{\eta}}_2]_j|$. However, recall from our assumptions on initialization that $[\bar{\boldsymbol{\eta}}_1]_j^2 - [\bar{\boldsymbol{\eta}}_2]_j^2 \geq \lambda > 0$, so (75) can only hold with strict inequality. Therefore,

$$g(\nu) := \sum_{j=1}^m [\mathbf{s}]_j \left(\frac{[\bar{\boldsymbol{\eta}}_1]_j^2 + [\bar{\boldsymbol{\eta}}_2]_j^2}{2} \sinh(2[\mathbf{s}]_j \nu) + [\bar{\boldsymbol{\eta}}_1]_j [\bar{\boldsymbol{\eta}}_2]_j \cosh(2[\mathbf{s}]_j \nu) \right)$$

is a strictly increasing (hence invertible) function because it is a sum of m strictly increasing functions. Using this $g(\nu)$, (74) can be written as $\alpha^2 g(\nu) = y$, and by using the inverse of g , we have

$$\nu = -\int_0^\infty r(\tau) d\tau = g^{-1} \left(\frac{y}{\alpha^2} \right). \tag{76}$$

Plugging (76) into (73), we get

$$\begin{aligned}
 & \lim_{t \rightarrow \infty} \mathbf{v}_1(t) \\
 &= \mathbf{U}_1 \lim_{t \rightarrow \infty} \boldsymbol{\eta}_1(t) + \alpha(\mathbf{I}_{k_1} - \mathbf{U}_1 \mathbf{U}_1^T) \bar{\mathbf{v}}_1 \\
 &= \alpha \mathbf{U}_1 \left(\bar{\boldsymbol{\eta}}_1 \odot \cosh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) + \bar{\boldsymbol{\eta}}_2 \odot \sinh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) \right) + \alpha(\mathbf{I}_{k_1} - \mathbf{U}_1 \mathbf{U}_1^T) \bar{\mathbf{v}}_1, \\
 & \lim_{t \rightarrow \infty} \mathbf{v}_2(t) \\
 &= \mathbf{U}_2 \lim_{t \rightarrow \infty} \boldsymbol{\eta}_2(t) + \alpha(\mathbf{I}_{k_2} - \mathbf{U}_2 \mathbf{U}_2^T) \bar{\mathbf{v}}_2 \\
 &= \alpha \mathbf{U}_2 \left(\bar{\boldsymbol{\eta}}_1 \odot \sinh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) + \bar{\boldsymbol{\eta}}_2 \odot \cosh \left(g^{-1} \left(\frac{y}{\alpha^2} \right) \mathbf{s} \right) \right) + \alpha(\mathbf{I}_{k_2} - \mathbf{U}_2 \mathbf{U}_2^T) \bar{\mathbf{v}}_2.
 \end{aligned}$$

This finishes the proof.

Appendix L. Proof of Theorem 7

L.1. Convergence of loss to zero

We first show that given the conditions on initialization, the training loss $\mathcal{L}(\Theta(t))$ converges to zero. Recall from (10) that the linear fully-connected network can be written as

$$f_{\text{fc}}(\mathbf{x}; \Theta_{\text{fc}}) = \mathbf{x}^T \mathbf{W}_1 \mathbf{W}_2 \cdots \mathbf{W}_{L-1} \mathbf{w}_L.$$

From the definition of the training loss \mathcal{L} , it is straightforward to check that the gradient flow dynamics read

$$\begin{aligned} \dot{\mathbf{W}}_l &= -\nabla_{\mathbf{W}_l} \mathcal{L}(\Theta_{\text{fc}}) = -\mathbf{W}_{l-1}^T \cdots \mathbf{W}_1^T \mathbf{X}^T \mathbf{r} \mathbf{w}_L^T \mathbf{W}_{L-1}^T \cdots \mathbf{W}_{l+1}^T \text{ for } l \in [L-1], \\ \dot{\mathbf{w}}_L &= -\nabla_{\mathbf{w}_L} \mathcal{L}(\Theta_{\text{fc}}) = -\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T \mathbf{X}^T \mathbf{r}, \\ \mathbf{W}_l(0) &= \alpha \bar{\mathbf{W}}_l \text{ for } l \in [L-1], \\ \mathbf{w}_L(0) &= \alpha \bar{\mathbf{w}}_L, \end{aligned} \tag{77}$$

where $\mathbf{r} \in \mathbb{R}^n$ is the residual vector satisfying $[\mathbf{r}]_i = f_{\text{fc}}(\mathbf{x}_i; \Theta_{\text{fc}}) - y_i$, as defined in Section 2.1. From (77), we have

$$\begin{aligned} \mathbf{W}_l^T \dot{\mathbf{W}}_l &= \dot{\mathbf{W}}_{l+1}^T \mathbf{W}_{l+1}^T = -\mathbf{W}_l^T \cdots \mathbf{W}_1^T \mathbf{X}^T \mathbf{r} \mathbf{w}_L^T \mathbf{W}_{L-1}^T \cdots \mathbf{W}_{l+1}^T, \\ \dot{\mathbf{W}}_l^T \mathbf{W}_l &= \mathbf{W}_{l+1}^T \dot{\mathbf{W}}_{l+1}^T = -\mathbf{W}_{l+1}^T \cdots \mathbf{W}_{L-1}^T \mathbf{w}_L \mathbf{r}^T \mathbf{X} \mathbf{W}_1 \cdots \mathbf{W}_l, \end{aligned}$$

for any $l \in [L-2]$. From this, we have

$$\frac{d}{dt} \mathbf{W}_l^T \mathbf{W}_l = \frac{d}{dt} \mathbf{W}_{l+1}^T \mathbf{W}_{l+1},$$

and thus

$$\begin{aligned} \mathbf{W}_l(t)^T \mathbf{W}_l(t) - \mathbf{W}_{l+1}(t)^T \mathbf{W}_{l+1}(t) &= \mathbf{W}_l(0)^T \mathbf{W}_l(0) - \mathbf{W}_{l+1}(0)^T \mathbf{W}_{l+1}(0) \\ &= \alpha^2 \bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \alpha^2 \bar{\mathbf{W}}_{l+1}^T \bar{\mathbf{W}}_{l+1}, \end{aligned} \tag{78}$$

for any $l \in [L-2]$. Similarly, we have

$$\begin{aligned} \mathbf{W}_{L-1}(t)^T \mathbf{W}_{L-1}(t) - \mathbf{w}_L(t)^T \mathbf{w}_L(t) &= \mathbf{W}_{L-1}(0)^T \mathbf{W}_{L-1}(0) - \mathbf{w}_L(0)^T \mathbf{w}_L(0) \\ &= \alpha^2 \bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \alpha^2 \bar{\mathbf{w}}_L^T \bar{\mathbf{w}}_L. \end{aligned} \tag{79}$$

Let us now consider the time derivative of $\mathcal{L}(\Theta_{\text{fc}}(t))$. We have the following chain of upper bounds on the time derivative:

$$\begin{aligned} \frac{d}{dt} \mathcal{L}(\Theta_{\text{fc}}(t)) &= \nabla_{\Theta_{\text{fc}}} \mathcal{L}(\Theta_{\text{fc}}(t))^T \dot{\Theta}_{\text{fc}}(t) = -\|\nabla_{\Theta_{\text{fc}}} \mathcal{L}(\Theta_{\text{fc}}(t))\|_2^2 \\ &\leq -\|\nabla_{\mathbf{w}_L} \mathcal{L}(\Theta_{\text{fc}}(t))\|_2^2 = -\|\dot{\mathbf{w}}_L(t)\|_2^2 \\ &= -\|\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T \mathbf{X}^T \mathbf{r}\|_2^2. \end{aligned} \tag{80}$$

Note from (80) that if $\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T$ is full-rank, its minimum singular value is positive, and one can bound

$$\|\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T \mathbf{X}^T \mathbf{r}\|_2 \geq \sigma_{\min}(\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T) \|\mathbf{X}^T \mathbf{r}\|_2. \tag{81}$$

We now prove that the matrix $\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T$ is full-rank, and its minimum singular value is bounded from below by $\alpha^{L-1} \lambda^{(L-1)/2}$ for any $t \geq 0$. To show this, it suffices to show that

$$\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T \mathbf{W}_1 \cdots \mathbf{W}_{L-1} \succeq \alpha^{2L-2} \lambda^{L-1} \mathbf{I}_d. \quad (82)$$

Now,

$$\begin{aligned} & \mathbf{W}_{L-1}^T \cdots \mathbf{W}_2^T \mathbf{W}_1^T \mathbf{W}_1 \mathbf{W}_2 \cdots \mathbf{W}_{L-1} \\ \stackrel{(a)}{=} & \mathbf{W}_{L-1}^T \cdots \mathbf{W}_2^T (\mathbf{W}_2 \mathbf{W}_2^T + \alpha^2 \bar{\mathbf{W}}_1^T \bar{\mathbf{W}}_1 - \alpha^2 \bar{\mathbf{W}}_2 \bar{\mathbf{W}}_2^T) \mathbf{W}_2 \cdots \mathbf{W}_{L-1} \\ \stackrel{(b)}{\succeq} & \mathbf{W}_{L-1}^T \cdots \mathbf{W}_3^T \mathbf{W}_2^T \mathbf{W}_2 \mathbf{W}_2^T \mathbf{W}_2 \mathbf{W}_3 \cdots \mathbf{W}_{L-1} \\ \stackrel{(a)}{=} & \mathbf{W}_{L-1}^T \cdots \mathbf{W}_3^T (\mathbf{W}_3 \mathbf{W}_3^T + \alpha^2 \bar{\mathbf{W}}_2^T \bar{\mathbf{W}}_2 - \alpha^2 \bar{\mathbf{W}}_3 \bar{\mathbf{W}}_3^T) \mathbf{W}_3 \cdots \mathbf{W}_{L-1} \\ \stackrel{(b)}{\succeq} & \mathbf{W}_{L-1}^T \cdots \mathbf{W}_3^T (\mathbf{W}_3 \mathbf{W}_3^T)^2 \mathbf{W}_3 \cdots \mathbf{W}_{L-1} \\ = & \cdots \succeq (\mathbf{W}_{L-1}^T \mathbf{W}_{L-1})^{L-1}, \end{aligned}$$

where equalities marked in (a) used (78), and inequalities marked in (b) used the initialization conditions $\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l \succeq \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T$. Next, it follows from (79) that

$$\begin{aligned} (\mathbf{W}_{L-1}^T \mathbf{W}_{L-1})^{L-1} &= (\mathbf{w}_L \mathbf{w}_L^T + \alpha^2 \bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \alpha^2 \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T)^{L-1} \\ &\succeq \alpha^{2L-2} (\bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T)^{L-1} \\ &\stackrel{(c)}{\succeq} \alpha^{2L-2} \lambda^{L-1} \mathbf{I}_d. \end{aligned}$$

where (c) used the assumption that $\bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T \succeq \lambda \mathbf{I}_d$. This proves (82). Applying (82) to (80) then gives

$$\begin{aligned} \frac{d}{dt} \mathcal{L}(\Theta_{\text{fc}}(t)) &\leq -\|\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T \mathbf{X}^T \mathbf{r}\|_2^2 \\ &\leq -\sigma_{\min}(\mathbf{W}_{L-1}^T \cdots \mathbf{W}_1^T)^2 \|\mathbf{X}^T \mathbf{r}\|_2^2 \\ &\leq -\alpha^{2L-2} \lambda^{L-1} \|\mathbf{X}^T \mathbf{r}\|_2^2 \\ &\stackrel{(d)}{\leq} -\alpha^{2L-2} \lambda^{L-1} \sigma_{\min}(\mathbf{X})^2 \|\mathbf{r}\|_2^2 \\ &= -\alpha^{2L-2} \lambda^{L-1} \sigma_{\min}(\mathbf{X})^2 \mathcal{L}(\Theta_{\text{fc}}(t)), \end{aligned}$$

where (d) used the fact that \mathbf{X}^T is a full column rank matrix to apply a bound similar to (81). From this, we get

$$\mathcal{L}(\Theta_{\text{fc}}(t)) \leq \mathcal{L}(\Theta_{\text{fc}}(0)) \exp(-\alpha^{2L-2} \lambda^{L-1} \sigma_{\min}(\mathbf{X})^2 t),$$

hence proving $\mathcal{L}(\Theta_{\text{fc}}(t)) \rightarrow 0$ as $t \rightarrow \infty$.

L.2. Characterizing the limit point: $\alpha \rightarrow 0$ case

Now, we move on to characterize the limit points of the gradient flow, for the ‘‘active regime’’ case $\alpha \rightarrow 0$. This part of the proof is motivated from the analysis in Ji and Telgarsky [18].

Let \mathbf{u}_l and \mathbf{v}_l be the top left and right singular vectors of \mathbf{W}_l , for $l \in [L-1]$. Note that since \mathbf{W}_l varies over time, the singular vectors and singular value also vary over time. Similarly, let s_l be the largest singular value of \mathbf{W}_l . We will show that the linear coefficients $\beta_{\text{fc}}(\Theta_{\text{fc}}) = \mathbf{W}_1 \cdots \mathbf{W}_{L-1} \mathbf{w}_L$ align with \mathbf{u}_1 as $\alpha \rightarrow 0$, and \mathbf{u}_1 is in the subspace of $\text{row}(\mathbf{X})$ in the limit $\alpha \rightarrow 0$, hence proving that $\beta_{\text{fc}}(\Theta_{\text{fc}})$ is the minimum ℓ_2 norm solution in the limit $\alpha \rightarrow 0$.

First, note from (78) and (79) that if we take trace of both sides, we get

$$\begin{aligned} \|\mathbf{W}_l\|_{\text{F}}^2 - \|\mathbf{W}_{l+1}\|_{\text{F}}^2 &= \alpha^2(\|\bar{\mathbf{W}}_l\|_{\text{F}}^2 - \|\bar{\mathbf{W}}_{l+1}\|_{\text{F}}^2) \quad \text{for } l \in [L-2], \\ \|\mathbf{W}_{L-1}\|_{\text{F}}^2 - \|\mathbf{w}_L\|_2^2 &= \alpha^2(\|\bar{\mathbf{W}}_{L-1}\|_{\text{F}}^2 - \|\bar{\mathbf{w}}_L\|_2^2). \end{aligned}$$

Summing the equations above for $l, l+1, \dots, L-1$, we get

$$\|\mathbf{W}_l\|_{\text{F}}^2 - \|\mathbf{w}_L\|_2^2 = \alpha^2(\|\bar{\mathbf{W}}_l\|_{\text{F}}^2 - \|\bar{\mathbf{w}}_L\|_2^2). \quad (83)$$

Next, consider the operator norms (i.e., the maximum singular values), denoted as $\|\cdot\|_2$, of the matrices.

$$\begin{aligned} \|\mathbf{W}_l\|_2^2 &\geq \mathbf{u}_{l+1}^T \mathbf{W}_l^T \mathbf{W}_l \mathbf{u}_{l+1} \\ &\stackrel{(e)}{=} \mathbf{u}_{l+1}^T \mathbf{W}_{l+1} \mathbf{W}_{l+1}^T \mathbf{u}_{l+1} + \alpha^2 \mathbf{u}_{l+1}^T (\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T) \mathbf{u}_{l+1} \\ &= \|\mathbf{W}_{l+1}\|_2^2 + \alpha^2 \mathbf{u}_{l+1}^T (\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T) \mathbf{u}_{l+1} \\ &\geq \|\mathbf{W}_{l+1}\|_2^2 - \alpha^2 \|\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T\|_2 \quad \text{for } l \in [L-2], \\ \|\mathbf{W}_{L-1}\|_2^2 &\geq \frac{\mathbf{w}_L}{\|\mathbf{w}_L\|_2} \mathbf{W}_{L-1}^T \mathbf{W}_{L-1} \frac{\mathbf{w}_L}{\|\mathbf{w}_L\|_2} \\ &\stackrel{(f)}{=} \frac{\mathbf{w}_L}{\|\mathbf{w}_L\|_2} \mathbf{w}_L \mathbf{w}_L^T \frac{\mathbf{w}_L}{\|\mathbf{w}_L\|_2} + \alpha^2 \frac{\mathbf{w}_L}{\|\mathbf{w}_L\|_2} (\bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T) \frac{\mathbf{w}_L}{\|\mathbf{w}_L\|_2} \\ &\geq \|\mathbf{w}_L\|_2^2 - \alpha^2 \|\bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T\|_2. \end{aligned}$$

where (e) used (78) and (f) used (79). Summing the inequalities gives

$$\|\mathbf{W}_l\|_2^2 \geq \|\mathbf{w}_L\|_2^2 - \alpha^2 \sum_{k=1}^{L-1} \|\bar{\mathbf{W}}_k^T \bar{\mathbf{W}}_k - \bar{\mathbf{W}}_{k+1} \bar{\mathbf{W}}_{k+1}^T\|_2. \quad (84)$$

From (83) and (84), we get a bound on the gap between the second powers of the Frobenius norm (or the ℓ_2 norm of singular values) and operator norm (or the maximum singular value s_l) of \mathbf{W}_l :

$$\|\mathbf{W}_l(t)\|_{\text{F}}^2 - \|\mathbf{W}_l(t)\|_2^2 \leq \alpha^2(\|\bar{\mathbf{W}}_l\|_{\text{F}}^2 - \|\bar{\mathbf{w}}_L\|_2^2) + \alpha^2 \sum_{k=l}^{L-1} \|\bar{\mathbf{W}}_k^T \bar{\mathbf{W}}_k - \bar{\mathbf{W}}_{k+1} \bar{\mathbf{W}}_{k+1}^T\|_2, \quad (85)$$

which holds for any $t \geq 0$. The gap (85) implies that each \mathbf{W}_l , for $l \in [L-1]$, can be written as

$$\mathbf{W}_l(t) = s_l(t) \mathbf{u}_l(t) \mathbf{v}_l(t)^T + O(\alpha^2). \quad (86)$$

Next, we show that the ‘‘adjacent’’ singular vectors \mathbf{v}_l and \mathbf{u}_{l+1} align with each other as $\alpha \rightarrow 0$. To this end, we will get lower and upper bounds for a quantity $\mathbf{v}_l^T \mathbf{W}_{l+1} \mathbf{W}_{l+1}^T \mathbf{v}_l$.

$$\mathbf{v}_l^T \mathbf{W}_{l+1} \mathbf{W}_{l+1}^T \mathbf{v}_l = \mathbf{v}_l^T \mathbf{W}_l^T \mathbf{W}_l \mathbf{v}_l - \alpha^2 \mathbf{v}_l^T \bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l \mathbf{v}_l + \alpha^2 \mathbf{v}_l^T \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T \mathbf{v}_l$$

$$\begin{aligned}
 &\geq \|\mathbf{W}_l\|_2^2 - \alpha^2 \|\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T\|_2 \\
 &= s_l^2 - \alpha^2 \|\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T\|_2, \tag{87}
 \end{aligned}$$

$$\begin{aligned}
 \mathbf{v}_l^T \mathbf{W}_{l+1} \mathbf{W}_{l+1}^T \mathbf{v}_l &= \mathbf{v}_l^T (s_{l+1}^2 \mathbf{u}_{l+1} \mathbf{u}_{l+1}^T + \mathbf{W}_{l+1} \mathbf{W}_{l+1}^T - s_{l+1}^2 \mathbf{u}_{l+1} \mathbf{u}_{l+1}^T) \mathbf{v}_l \\
 &= s_{l+1}^2 (\mathbf{v}_l^T \mathbf{u}_{l+1})^2 + \mathbf{v}_l^T (\mathbf{W}_{l+1} \mathbf{W}_{l+1}^T - s_{l+1}^2 \mathbf{u}_{l+1} \mathbf{u}_{l+1}^T) \mathbf{v}_l \\
 &\leq s_{l+1}^2 (\mathbf{v}_l^T \mathbf{u}_{l+1})^2 + \|\mathbf{W}_{l+1}\|_F^2 - \|\mathbf{W}_{l+1}\|_2^2. \tag{88}
 \end{aligned}$$

Combining (87), (88), and (85), we get

$$\begin{aligned}
 s_l^2 &\leq s_{l+1}^2 (\mathbf{v}_l^T \mathbf{u}_{l+1})^2 + \alpha^2 \|\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T\|_2 + \|\mathbf{W}_{l+1}\|_F^2 - \|\mathbf{W}_{l+1}\|_2^2 \\
 &\leq s_{l+1}^2 (\mathbf{v}_l^T \mathbf{u}_{l+1})^2 + \alpha^2 (\|\bar{\mathbf{W}}_{l+1}\|_F^2 - \|\bar{\mathbf{w}}_L\|_2^2) + \alpha^2 \sum_{k=l}^{L-1} \|\bar{\mathbf{W}}_k^T \bar{\mathbf{W}}_k - \bar{\mathbf{W}}_{k+1} \bar{\mathbf{W}}_{k+1}^T\|_2. \tag{89}
 \end{aligned}$$

Next, by a similar reasoning as (87), we have

$$s_l^2 \geq \mathbf{u}_{l+1}^T \mathbf{W}_l^T \mathbf{W}_l \mathbf{u}_{l+1} \geq s_{l+1}^2 - \alpha^2 \|\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T\|_2. \tag{90}$$

Combining (89) and (90) and dividing both sides by s_{l+1}^2 , we get

$$(\mathbf{v}_l(t)^T \mathbf{u}_{l+1}(t))^2 \geq 1 - \alpha^2 \frac{G_l}{s_{l+1}(t)^2} \tag{91}$$

for $t \geq 0$, where

$$G_l := \|\bar{\mathbf{W}}_l^T \bar{\mathbf{W}}_l - \bar{\mathbf{W}}_{l+1} \bar{\mathbf{W}}_{l+1}^T\|_2 + (\|\bar{\mathbf{W}}_{l+1}\|_F^2 - \|\bar{\mathbf{w}}_L\|_2^2) + \sum_{k=l}^{L-1} \|\bar{\mathbf{W}}_k^T \bar{\mathbf{W}}_k - \bar{\mathbf{W}}_{k+1} \bar{\mathbf{W}}_{k+1}^T\|_2.$$

By a similar argument, we can also get

$$\frac{(\mathbf{v}_{L-1}(t)^T \mathbf{w}_L(t))^2}{\|\mathbf{w}_L(t)\|_2^2} \geq 1 - \alpha^2 \frac{G_{L-1}}{\|\mathbf{w}_L(t)\|_2^2}, \tag{92}$$

where

$$G_{L-1} := 2 \|\bar{\mathbf{W}}_{L-1}^T \bar{\mathbf{W}}_{L-1} - \bar{\mathbf{w}}_L \bar{\mathbf{w}}_L^T\|_2.$$

From (91) and (92), we can note that as $\alpha \rightarrow 0$, the inner product between the adjacent singular vectors converges to ± 1 , unless $s_2, \dots, s_{L-1}, \|\mathbf{w}_L\|_2$ also diminish to zero. So it is left to show that the singular values do not diminish to zero as $\alpha \rightarrow 0$. To this end, recall that we proved in the previous subsection that

$$\lim_{t \rightarrow \infty} \mathbf{X} \mathbf{W}_1(t) \cdots \mathbf{W}_{L-1}(t) \mathbf{w}_L(t) = \mathbf{y}.$$

A necessary condition for this to hold is that

$$\frac{\|\mathbf{y}\|_2}{\|\mathbf{X}\|_2} \leq \lim_{t \rightarrow \infty} \|\mathbf{W}_1(t) \cdots \mathbf{W}_{L-1}(t) \mathbf{w}_L(t)\|_2 \leq \lim_{t \rightarrow \infty} \prod_{l=1}^{L-1} s_l(t) \|\mathbf{w}_L(t)\|_2.$$

This means that after converging to the global minimum solution of the problem (i.e., $t \rightarrow \infty$), the product of the singular values must be at least greater than some constant independent of α . Moreover, we can see from (87) and (90) that the gap between singular values squared of adjacent layers is bounded by $O(\alpha^2)$, for all $t \geq 0$; so the maximum singular values become closer and closer to each other as α diminishes. This implies that

$$\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} s_l(t) \geq \frac{\|\mathbf{y}\|_2^{1/L}}{\|\mathbf{X}\|_2^{1/L}} \quad \text{for } l \in [L-1], \quad \lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} \|\mathbf{w}_L(t)\|_2 \geq \frac{\|\mathbf{y}\|_2^{1/L}}{\|\mathbf{X}\|_2^{1/L}}.$$

Therefore, we have the alignment of singular vectors at convergence as $\alpha \rightarrow 0$:

$$\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} (\mathbf{v}_l(t)^T \mathbf{u}_{l+1}(t))^2 = 1, \quad \text{for } l \in [L-2], \quad \lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} \frac{(\mathbf{v}_{L-1}(t)^T \mathbf{w}_L(t))^2}{\|\mathbf{w}_L(t)\|_2^2} = 1. \quad (93)$$

So far, we saw from (86) that $\mathbf{W}_l(t)$'s become rank-1 matrices as $\alpha \rightarrow 0$, and from (93) that the top singular vectors align with each other as $t \rightarrow \infty$ and $\alpha \rightarrow 0$. These imply that, as $t \rightarrow \infty$ and $\alpha \rightarrow 0$, $\beta_{\text{fc}}(\Theta_{\text{fc}})$ is a scalar multiple of the \mathbf{u}_1 , the top left singular vector of \mathbf{W}_1 :

$$\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} \beta_{\text{fc}}(\Theta_{\text{fc}}(t)) = c \cdot \lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} \mathbf{u}_1(t), \quad (94)$$

for some $c \in \mathbb{R}$.

In light of (94), it remains to take a close look at $\mathbf{u}_1(t)$. Note from the gradient flow dynamics of \mathbf{W}_1 that $\dot{\mathbf{W}}_1$ is always a rank-1 matrix whose columns are in the row space of \mathbf{X} , since $\mathbf{X}^T \mathbf{r} \in \text{row}(\mathbf{X})$. This implies that, if we decompose \mathbf{W}_1 into two orthogonal components \mathbf{W}_1^\perp and \mathbf{W}_1^\parallel so that the columns in \mathbf{W}_1^\parallel are in $\text{row}(\mathbf{X})$ and the columns in \mathbf{W}_1^\perp are in the orthogonal subspace $\text{row}(\mathbf{X})^\perp$, we have

$$\dot{\mathbf{W}}_1^\perp = \mathbf{0}, \quad \dot{\mathbf{W}}_1^\parallel = \dot{\mathbf{W}}_1.$$

That is, any component $\mathbf{W}_1^\perp(0)$ orthogonal to $\text{row}(\mathbf{X})$ remains unchanged for all $t \geq 0$, while the component \mathbf{W}_1^\parallel changes by the gradient flow. Since we have

$$\left\| \mathbf{W}_1^\perp(t) \right\|_{\text{F}} = \left\| \mathbf{W}_1^\perp(0) \right\|_{\text{F}} \leq \alpha \left\| \bar{\mathbf{W}}_1 \right\|_{\text{F}},$$

the component in \mathbf{W}_1 that is orthogonal to $\text{row}(\mathbf{X})$ diminishes to zero as $\alpha \rightarrow 0$. This means that at the limit $\alpha \rightarrow 0$, the columns of \mathbf{W}_1 are entirely from $\text{row}(\mathbf{X})$, which also means that

$$\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow \infty} \beta_{\text{fc}}(\Theta_{\text{fc}}(t)) \in \text{row}(\mathbf{X}).$$

However, recall that there is only one unique global minimum of $\mathbf{X}\mathbf{z} = \mathbf{y}$ in $\text{row}(\mathbf{X})$: namely, $\mathbf{z} = \mathbf{X}^T(\mathbf{X}\mathbf{X}^T)^{-1}\mathbf{y}$, the minimum ℓ_2 norm solution. This finishes the proof.