

Depth, Not Data: An Analysis of Hessian Spectral Bifurcation

Shenyang Deng
Department of Computer Science
Dartmouth College
Hanover, NH, USA
shenyang.deng.gr@dartmouth.edu

Boyao Liao
Department of Mathematics
University of Birmingham
Birmingham, UK
bxl307@student.bham.ac.uk

Zhuoli Ouyang
Department of Computer Science
Dartmouth College
Hanover, NH, USA
Zhuoli.Ouyang@dartmouth.edu

Tianyu Pang
Department of Computer Science
Dartmouth College
Hanover, NH, USA
tianyu.pang.gr@dartmouth.edu

Yaoqing Yang
Department of Computer Science
Dartmouth College
Hanover, NH, USA
yaoqing.yang@dartmouth.edu

Abstract—The eigenvalue distribution of the Hessian matrix plays a crucial role in understanding the optimization landscape of deep neural networks. Prior work has attributed the well-documented “bulk-and-spike” spectral structure, where a few dominant eigenvalues are separated from a bulk of smaller ones, to the imbalance in the data covariance matrix. In this work, we challenge this view by demonstrating that such spectral *Bifurcation* can arise purely from the network architecture, independent of data imbalance.

Specifically, we analyze a deep linear network setup and prove that, even when the data covariance is perfectly balanced, the Hessian still exhibits a *Bifurcation* eigenvalue structure: a dominant cluster and a bulk cluster. Crucially, we establish that the ratio between dominant and bulk eigenvalues scales linearly with the network depth. This reveals that the spectral gap is strongly affected by the network architecture rather than solely by data distribution. Our results suggest that both model architecture and data characteristics should be considered when designing optimization algorithms for deep networks.

Index Terms—Optimization, Deep Learning Theory, Hessian Spectrum Analysis

I. INTRODUCTION

As the application scope of Deep Neural Networks (DNNs) continues to expand, there is a pressing need to improve optimization algorithms designed for these models. Achieving such algorithmic improvements requires a deeper theoretical understanding of the loss landscape of DNNs because the geometry of the loss landscape often strongly affects the effectiveness of optimization [1]–[4] and the quality of the converged model (e.g., generalization performance) [5]–[7].

In this context, the Hessian matrix has become a critical tool to characterize the curvature of the loss landscape. Extensive empirical research has demonstrated that the loss landscapes of various neural networks exhibit an ill-conditioned eigenvalue distribution structure [8]–[13]. Specifically: [8] and [9] observed that the gradient descent dynamics typically happen in a “tiny subspace” spanned by a few dominant Hessian eigenvectors, while the majority of directions are flat. [10] and [11]

provided a detailed spectral analysis, revealing that the Hessian spectrum is composed of massive “bulk” concentrated near zero and a few “outlier” eigenvalues (spikes). Furthermore, [14] and [15] argued that these spectral distributions are often heavy-tailed, suggesting a form of implicit self-regularization.

Complementing these empirical phenomena, theoretical work has offered explanations in specific settings. For example, [16] and [17] use Random Matrix Theory (RMT) to build a direct link between the Hessian spectrum and the covariance matrix of the input data. A prevailing conclusion in these works [18], [19] is that the eigenvalue distribution of the Hessian $H_L = \nabla^2 L(W)$ is heavily entangled with the distribution of the data covariance matrix. Consequently, the common perspective is that the “imbalance” observed in the Hessian spectrum—where a few large eigenvalues dominate—is primarily inherited from the intrinsic imbalance of the data covariance matrix.

This brings us to the core of our work. Following the context established by the aforementioned bulk-plus-spike literature, we classify the eigenvalues into a “bulk space” and a “dominant spike space.” Let $\lambda_1, \dots, \lambda_p$ be the non-zero eigenvalues of the Hessian H_L . We typically observe an index $k < p$ that separates the eigenvalues into two distinct clusters with a significant gap:

$$\underbrace{\lambda_1, \dots, \lambda_k}_{\text{Dominant Spikes}} \in [ma, m(a + \delta)], \quad \underbrace{\lambda_{k+1}, \dots, \lambda_p}_{\text{Bulk}} \in [a, a + \delta], \quad (1)$$

where $\delta \in (0, a)$ represents a fluctuation amplitude of eigenvalues in the corresponding eigensubspace, and $m \gg 1$ represents the magnitude of the spectral gap. In this work, we do not consider zero eigenvalues, as their corresponding eigenspaces contain no meaningful information. While prior literature typically attributes this sharp **spectral bifurcation**—where the top k eigenvalues separate from the bulk by a large magnitude m —to the imbalance of the data distribution, our

work challenges this prevailing view by posing the following question:

Key Question

Q1: Is the spectral bifurcation of H_L (with k spikes and gap m) exclusively attributable to the spectral imbalance of the input data covariance?

Our work demonstrates that an imbalance in the data covariance (or cross-covariance between inputs and outputs) spectrum is not a necessary condition for such spectral bifurcation. Although we acknowledge that the Hessian in DNN optimization is indeed influenced by imbalanced data, we show that even when the data is perfectly balanced (i.e., the data is whitened such that the covariance and cross covariance are identity-like or uniform), the Hessian of the neural network optimization problem still generates a bifurcation spectral structure.

We show in a theoretical setting that there exists a k such that the top k eigenvalues are orders of magnitude (m times) larger than the remaining $p - k$ non-zero eigenvalues. Crucially, the severity of this ill-conditioning (represented by m) is intrinsic to the model design and correlates with the network depth/architecture, rather than being solely dependent on data distribution.

To give a clearer intuition, we first visualize this conclusion via a numerical simulation, as shown in Figure 1. We train a deep linear neural network to learn a simple linear operator mapping ($\mathbf{y} = \Phi \mathbf{x}$). We utilize a sufficient quantity of samples to ensure the data is whitened such that the cross-correlation matrix satisfies $\Sigma_{yx} = \mathbb{E}[\mathbf{y}\mathbf{x}^\top] = \Phi \simeq I_r$ (similar to the setting in [20], where I_r is an $r \times r$ identity matrix, \simeq is the 0 padding operator; the definition can be found in Section II). As illustrated, even when the eigenvalues of the data distribution are uniform, the Hessian of the network exhibits a clear stratification into two non-zero eigenvalue clusters and a cluster near zero.

Following the theoretical settings of [20] and [21] for deep linear networks, we provide a theoretical answer to this problem. We analyze why this imbalanced structure forms and identify the factors influencing it. Under the deep linear network setup following [20] and [21], our theoretical contributions are summarized as follows:

- **C1 (Spectral Bifurcation Independence):** We prove that even when both data covariance and cross-covariance are perfectly balanced (i.e., $\Sigma_{xx} \simeq I_{d_*}$ and $\Sigma_{yx} = UV^\top$ where U, V are column-orthogonal matrices), the Hessian still exhibits a *Bifurcation* structure among its non-zero eigenvalues: a dominant cluster and a bulk cluster. The remaining eigenvalues are zero and correspond to directions outside the data support. This shows that data imbalance is *not* a necessary condition for spectral bifurcation.
- **C2 (Depth-Dependent Gap):** We establish that the spectral gap between the dominant and bulk clusters scales linearly with the network depth L . Specifically, the

dominant eigenvalues are approximately L times larger than the bulk eigenvalues, i.e., $m = \Theta(L)$ as illustrated in Figure 1. This reveals that the ill-conditioning severity is intrinsically determined by the model architecture rather than solely by data distribution.

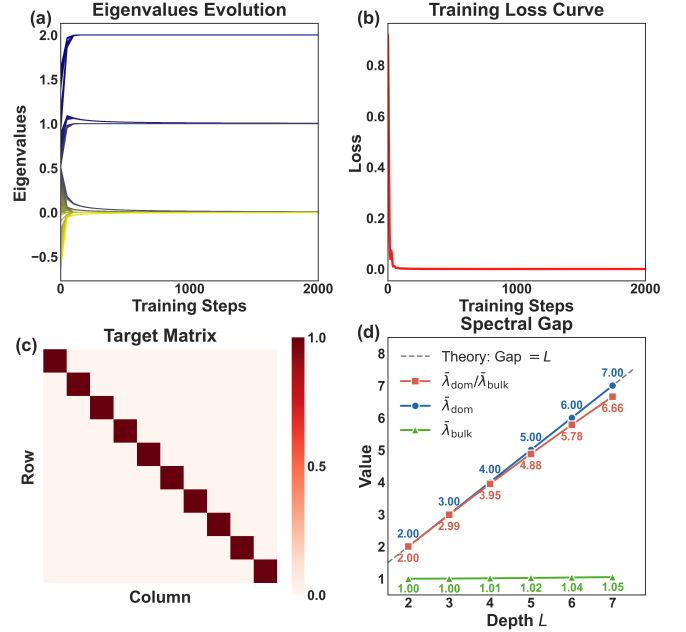


Fig. 1: (a) Hessian eigenvalue evolution for $L = 2$ on whitened data. (b) Training loss curve ($L = 2$). (c) Target matrix (whitened data). (d) Spectral gap $\bar{\lambda}_{\text{dom}}/\bar{\lambda}_{\text{bulk}}$ across depths $L = 2-7$, where $\bar{\lambda}_{\text{dom}}$ and $\bar{\lambda}_{\text{bulk}}$ denote the mean eigenvalues of the dominant and bulk spaces, respectively; the empirical gap closely matches the theoretical prediction $\text{Gap} = L$.

II. PRELIMINARY

Notation. For matrix products, we define $W^{k:l} = W^k \dots W^l$ when $k > l$ and $W^{k:l} = (W^k)^\top \dots (W^l)^\top$ when $k < l$. Additional operations include: \otimes for Kronecker product, vec_r for row-wise matrix vectorization, $\text{rk}(W)$ for matrix rank, and I_k for the $k \times k$ identity matrix. Let matrix $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{p \times q}$. To work with matrix-matrix derivatives in the row-wise vectorization form, we adopt the convention:

$$\frac{\partial A}{\partial B} := \frac{\partial \text{vec}_r(A)}{\partial \text{vec}_r(B)^\top}.$$

A useful identity for our analysis is:

$$\frac{\partial (AWB)}{\partial W} = A \otimes B^\top.$$

If $m \geq p$ and $n \geq q$, we define the padding operator \simeq as:

$$A \simeq B \Leftrightarrow A = \begin{bmatrix} B & \mathbf{0}_{p \times (n-q)} \\ \mathbf{0}_{(m-p) \times q} & \mathbf{0}_{(m-p) \times (n-q)} \end{bmatrix},$$

where $\mathbf{0}_{i \times j}$ denotes an $i \times j$ zero matrix.

Deep Linear Neural Network. Following the setup of [20]–[22], we consider a depth- L deep linear neural network defined by the composition $F(\mathbf{x}) = W^L \dots W^1 \mathbf{x} = W^{L:1} \mathbf{x}$, where

the weight matrix $W^l \in \mathbb{R}^{d_l \times d_{l-1}}$ parameterizes the l -th layer for $l = 1, \dots, L$. The network has input dimension $d = d_0$, output dimension $K = d_L$, hidden layer widths d_1, \dots, d_{L-1} , and total number of hidden neurons $D = \sum_{l=1}^{L-1} d_l$. We denote the total parameter set as $W := \{W^1, \dots, W^L\}$ and use the explicit parameterization $F_W(\mathbf{x}) = W^{L:1}\mathbf{x}$ when necessary.

Loss Function. We consider a training dataset of i.i.d. samples $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ from distribution p . We consider the squared loss $L_W(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2$ where $\hat{\mathbf{y}} = F_W(\mathbf{x})$. The population loss is defined as $L(W) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p} [L_W(\mathbf{x}, \mathbf{y})]$. For brevity, specifically in subsequent derivations, we denote the expectation over the data distribution simply as \mathbb{E} . The Hessian of the population loss is denoted $H = \nabla^2 L(W)$. We define the residual at sample (\mathbf{x}, \mathbf{y}) as $\delta_{\mathbf{x}, \mathbf{y}} := \hat{\mathbf{y}} - \mathbf{y}$.

Applying the above to the network output, we obtain:

$$\frac{\partial F}{\partial W^l} = W^{L:l+1} \otimes (\mathbf{x}^\top W^{1:l-1}) \in \mathbb{R}^{d_L \times d_l d_{l-1}}.$$

Using the chain rule (backpropagation), the gradient of the loss for a sample (\mathbf{x}, \mathbf{y}) with respect to W^l is:

$$\begin{aligned} \nabla_{W^l} L(W_t) &= [W^{l+1:L} \delta_{\mathbf{x}, \mathbf{y}}] [W^{l-1:1} \mathbf{x}]^\top \\ &= W^{l+1:L} (W^{L:1} \mathbf{x} \mathbf{x}^\top - \mathbf{y} \mathbf{x}^\top) W^{1:l-1}. \end{aligned}$$

For subsequent analysis, we introduce the following shorthand notations for the expected quantities (where \mathbb{E} denotes $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu}$):

$$\begin{aligned} \Omega &:= \mathbb{E}[\delta_{\mathbf{x}, \mathbf{y}} \mathbf{x}^\top] = \mathbb{E}[W^{L:1} \mathbf{x} \mathbf{x}^\top] - \mathbb{E}[\mathbf{y} \mathbf{x}^\top], \\ \Sigma_{xx} &:= \mathbb{E}[\mathbf{x} \mathbf{x}^\top], \quad \Sigma_{yx} := \mathbb{E}[\mathbf{y} \mathbf{x}^\top]. \end{aligned}$$

Gauss-Newton Decomposition. Following the analytical technique in [21], [23], we apply the Gauss-Newton Decomposition to analyze the Hessian of a deep linear neural network. This technique involves examining the Hessian through the lens of the chain rule, decomposing the total loss Hessian $H_L = H_o + H_f$, where:

$$H_o = \mathbb{E} \left[\nabla_W F(\mathbf{x})^\top \left[\frac{\partial^2 L_W}{\partial F^2} \right] \nabla_W F(\mathbf{x}) \right],$$

and

$$H_f = \mathbb{E} \left[\sum_{c=1}^{d_L} \left[\frac{\partial L_W}{\partial F_c} \right] \nabla_W^2 F_c(\mathbf{x}) \right].$$

In the case of Mean Squared Error (MSE) loss, this Gauss-Newton breakdown aligns precisely with the earlier analysis as shown in [21]. Following convention, we call the initial component H_o the *outer-product Hessian* and the latter component H_f the *functional Hessian*.

Balanced Initialization. We follow the balanced initialization procedure in [20]:

Initialization II.1. Let $d_0, d_1, \dots, d_L \in \mathbb{N}$ satisfy $\min\{d_1, \dots, d_{L-1}\} \geq \min\{d_0, d_L\}$, and let \mathcal{D} be a distribution over $d_L \times d_0$ matrices. A *balanced initialization* of the weight matrices $W^l \in \mathbb{R}^{d_l \times d_{l-1}}$, for $l = 1, \dots, L$, is

performed as follows:

- 1) Sample a matrix $A \in \mathbb{R}^{d_L \times d_0}$ according to \mathcal{D} .
- 2) Compute the singular value decomposition $A = U \Sigma_0 V^\top$, where $U \in \mathbb{R}^{d_L \times \min\{d_0, d_L\}}$ and $V \in \mathbb{R}^{d_0 \times \min\{d_0, d_L\}}$ have orthonormal columns, and $\Sigma_0 \in \mathbb{R}^{\min\{d_0, d_L\} \times \min\{d_0, d_L\}}$ is diagonal, containing the singular values of A .
- 3) Set the weight matrices as $W^L \simeq U \Sigma_0^{1/L}$, $W^{L-1} \simeq \Sigma_0^{1/L}$, \dots , $W^2 \simeq \Sigma_0^{1/L}$, $W^1 \simeq \Sigma_0^{1/L} V^\top$.

Gradient Descent Dynamics. We optimize the network parameters $W = \{W^1, \dots, W^L\}$ using Gradient Descent (GD) with a constant learning rate $\eta > 0$. Following [20], we assume the data is whitened and consider the gradient update in the expectation sense. Substituting the gradient derivation from the previous section, the population gradient update takes the explicit form:

$$W_{t+1}^l = W_t^l - \eta \cdot W_t^{l+1:L} (W_t^{L:1} \Sigma_{xx} - \Sigma_{yx}) W_t^{1:l-1} \quad (2)$$

III. HESSIAN BIFURCATION THEORY

A. Theorem Assumptions

We have the following assumptions for our results. However, note that not all assumptions are actually used in the theorem. Some are included solely to help the reader develop intuition for our results by considering a more concrete special case.

Assumption III.1 (Sufficient Width & Whitened Input). We impose structural constraints on the network width to prevent information propagation bottlenecks, assuming $d_* = \min\{d_L, d_0\} \leq \min\{d_{L-1}, \dots, d_1\}$. Furthermore, regarding the data distribution, we assume the input covariance matrix is whitened (isotropic):

$$\Sigma_{xx} \simeq I_{d_*}, \quad (3)$$

where I_{d_*} is the $d_* \times d_*$ identity matrix.

Assumption III.2 (Alignment with Initialization). We assume the target mapping is aligned with the principal directions of A in **Initialization II.1**. Let U and V be the singular vector matrices defined in the initialization procedure. To align with the data rank r , the cross-covariance matrix satisfies:

$$\Sigma_{yx} = U \mathcal{I}_r V^\top, \quad (4)$$

where $\mathcal{I}_r \in \mathbb{R}^{d_* \times d_*} \simeq I_r$.

Assumption III.3 (Uniform Spectral Initialization (USI)). This is a special case that we will consider later. We assume the singular values of the initial weight product are uniform. Denoting the diagonal singular value matrix as Σ in **Initialization II.1** (corresponding to Σ_0 in the initialization definition), we have:

$$\Sigma_0^{1/L} \simeq \mu I_r, \quad (5)$$

where $\mu > 0$ is a constant. This ensures network operates as a scaled isometry on the full parameter space at initialization.

B. Main Result

To establish our main result on the spectral structure of the Hessian, we first need to understand how the weight matrices evolve during training. A key observation is that all weight matrices in a deep linear network share a common spectral structure throughout the optimization process. This property, which we formalize in the following lemma III.4, serves as the foundation for our subsequent analysis.

Lemma III.4 (Shared Spectral Structure). *Under Assumptions III.1 and III.2, all weight matrices W_t^k for $k = 1, \dots, L$ share a same spectral structure $\Sigma_t^{1/L} = \text{diag}(\lambda_{1,t}, \lambda_{2,t}, \dots, \lambda_{d_*,t})$ at any time $t \geq 0$. Specifically:*

- 1) For $1 < k < L$: $W_t^k \simeq \Sigma_t^{1/L}$.
- 2) For $k = 1$: $W_t^1 \simeq \Sigma_t^{1/L} V^\top$.
- 3) For $k = L$: $W_t^L \simeq U \Sigma_t^{1/L}$.

where $U \in \mathbb{R}^{d_L \times d_*}$ and $V \in \mathbb{R}^{d_0 \times d_*}$ are the left and right singular vector matrices from the balanced initialization, which remain constant throughout training. Furthermore, the eigenvalues $\lambda_{i,t}$ for $i = 1, \dots, r$ evolve according to:

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1}, \quad (6)$$

while the eigenvalues $\lambda_{i,t}$ for $i = r+1, \dots, d_*$ remain unchanged at their initial values.

Explain. Lemma III.4 reveals that under **Initialization II.1**, the spectral structure of all weight matrices is governed by a single diagonal matrix $\Sigma_t^{1/L}$, whose diagonal entries $\lambda_{i,t}$ evolve according to a unified update rule. This shared spectral structure is crucial because it implies that the optimization dynamics can be characterized entirely by tracking the evolution of these d_* eigenvalues.

Building upon this spectral characterization, we are now ready to present our main theorem, which establishes the bifurcation structure of the Hessian eigenvalues. The theorem demonstrates that even when the data covariance is perfectly balanced (i.e., $\Sigma_{xx} \simeq I_{d_*}$), the Hessian of the loss function exhibits a pronounced two-cluster spectral structure, with a gap that scales linearly with the network depth L .

Theorem III.5 (Hessian Bifurcation). *Under Assumptions III.1 and III.2, and assuming $r \leq d_*$, consider a depth- L deep linear neural network trained with gradient descent with step size $\eta < \min\{\frac{1}{L}, \frac{1}{M^{2L-2}}\}$. Let $\lambda_{i,t}$ for $i = 1, \dots, r$ denote the effective eigenvalues of the weight matrix $\Sigma_t^{1/L}$ at time t . Suppose that $\lambda_{i,t} \in [m_t - \delta_t, m_t + \delta_t]$ and the condition $(m_t + \delta_t)/(m_t - \delta_t) < L^{2/(L-1)}$ holds. Then, the Hessian $H_{L,t}$ exhibits a two-cluster spectral structure:*

- 1) **(Dominant Space)** There exist r^2 eigenvalues of $H_{L,t}$ lying in:

$$\left[\begin{array}{l} L(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t}), \\ L(m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t}) \end{array} \right]. \quad (7)$$

- 2) **(Bulk Space)** There exist $(d_* + d_L - 2r)r$ eigenvalues of $H_{L,t}$ lying in:

$$\left[\begin{array}{l} (m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t}), \\ (m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t}) \end{array} \right]. \quad (8)$$

- 3) **(Zero Space)** The remaining eigenvalues of $H_{L,t}$ are zero.

Moreover, let λ_{dom} denote an arbitrary eigenvalue belonging to the Dominant Space and λ_{bulk} denote an arbitrary eigenvalue belonging to the Bulk Space. Their ratio satisfies:

$$\frac{\lambda_{dom}}{\lambda_{bulk}} = \Theta(L), \quad (9)$$

where α is a positive constant independent of t , provided t is sufficiently large.

Explain. Theorem III.5 provides a complete characterization of the Hessian spectrum in deep linear networks. The result reveals three distinct eigenspaces: (i) a dominant space containing r^2 eigenvalues of order $L \cdot m_t^{2(L-1)}$, (ii) a bulk space containing $(d_* + d_L - 2r)r$ eigenvalues of order $m_t^{2(L-1)}$, and (iii) a zero space containing the remaining eigenvalues that vanish asymptotically. Crucially, the ratio between the dominant and bulk eigenvalues is $\Theta(L)$, demonstrating that the spectral gap is an intrinsic property of the network architecture rather than a consequence of data imbalance.

Remark. The condition $(m_t + \delta_t)/(m_t - \delta_t) < L^{1/(2(L-1))}$ ensures that the eigenvalues remain sufficiently concentrated to maintain a clear separation between the two clusters. This condition is naturally satisfied when the initialization is near-uniform and becomes increasingly easier to satisfy as the network depth L grows.

To gain deeper insight into the mechanism underlying this spectral bifurcation, we consider a special case where the weight matrices are initialized with uniform singular values. This setting, formalized in Assumption III.3, allows us to obtain a sharper characterization of the Hessian spectrum.

Corollary III.6 (Hessian Bifurcation with USI). *Under Assumptions III.1, III.2, and III.3, and assuming $r \leq d_*$, the Hessian $H_{L,t}$ has exactly two distinct nonzero eigenvalues (up to an exponentially small perturbation). Specifically:*

- 1) **(Dominant Space)** There exist r^2 eigenvalues equal to $L\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})$.
- 2) **(Bulk Space)** There exist $(d_* + d_L - 2r)r$ eigenvalues equal to $\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})$.
- 3) **(Zero Space)** The remaining eigenvalues are $O(e^{-L\alpha\eta t})$.

Moreover, the ratio between the dominant and bulk eigenvalues is exactly L , i.e.,

$$\frac{\lambda_{dom}}{\lambda_{bulk}} = L, \quad (10)$$

where μ_t denotes the common value of all effective eigenvalues $\lambda_{i,t} = \mu_t$ for $i = 1, \dots, r$ at time t , and α is a positive constant.

Corollary III.6 reveals the fundamental mechanism behind the Hessian bifurcation phenomenon. Under uniform spectral initialization, the Hessian possesses exactly two distinct nonzero eigenvalues: $L\mu_t^{2(L-1)}$ in the dominant space and $\mu_t^{2(L-1)}$ in the bulk space, with a ratio of exactly L . This clean separation provides a transparent view of how the spectral structure emerges.

Proof Sketch. For detailed proofs of the lemma and theorems, you can refer to the full version of this paper on arXiv [24]. Under Assumptions III.1 and III.2, we first establish that all weight matrices share a common spectral structure throughout training (Lemma III.4), where the singular values of the end-to-end mapping evolve according to a unified dynamical system $\Sigma_{t+1}^{1/L} = g(\Sigma_t^{1/L})$. Analyzing this dynamics shows that the population loss converges exponentially as $L(W_t) = O(e^{-2L\alpha\eta t})$, which by the Gauss-Newton decomposition $H_{L,t} = H_{o,t} + H_{f,t}$ implies that the functional Hessian norm satisfies $\|H_{f,t}\|_2 = O(e^{-L\alpha\eta t})$ (Lemma B.8). Following [21] Proposition 2, the outer-product Hessian admits the factorization $H_{o,t} = A_{o,t}B_oA_{o,t}^\top$, whose nonzero eigenvalues coincide with those of the Gram matrix $B_o^{1/2}A_{o,t}^\top A_{o,t}B_o^{1/2}$. By explicitly computing the weight matrix products and exploiting the Kronecker product structure, we show that these eigenvalues split into a dominant cluster of r^2 eigenvalues scaling as $L \cdot m_t^{2(L-1)}$ and a bulk cluster of $(d_* + d_L - 2r)r$ eigenvalues scaling as $m_t^{2(L-1)}$ (Lemma B.9). Finally, applying Weyl’s inequality to $H_{L,t} = H_{o,t} + H_{f,t}$ shows that the eigenvalues of the true Hessian $H_{L,t}$ lie within $O(e^{-L\alpha\eta t})$ of those of $H_{o,t}$, yielding the stated two-cluster structure with ratio $\lambda_{\text{dom}}/\lambda_{\text{bulk}} = \Theta(L)$.

IV. SIMULATION

To illustrate the theoretical results in Sections II and III, we conduct numerical simulations under Assumptions III.1, III.2. For brevity, we provide 2 experiments here, and detailed experimental settings and additional results are provided in Appendix C.

We consider a depth- L deep linear network with input dimension $d_0 = 10$, output dimension $d_L = 16$, hidden widths $d_1 = \dots = d_{L-1} = 20$, and effective rank $r = 4$, trained on whitened data ($\Sigma_{xx} \simeq I_{d_*}$) with $L \in \{2, 3\}$. Figure 2 reports the Hessian eigenvalue trajectories and training loss curves, the latter confirming convergence. The non-zero spectrum cleanly separates into a dominant cluster of dimension $r^2 = 16$ and a bulk cluster of dimension $(d_* + d_L - 2r)r = 72$, consistent across both depths. Moreover, the spectral gap scales with depth, yielding ratios of approximately 2 for $L = 2$ and 3 for $L = 3$, matching the linear growth of dominant eigenvalues with L predicted by Theorem III.5.

Nonlinear Discussion. To probe whether the Hessian bifurcation extends beyond the linear regime, we repeat the experiment with a tanh activation applied after each layer, keeping all other settings unchanged. As shown in Figure 3, the non-zero spectrum still splits cleanly into a dominant and a bulk cluster (panel a), and training converges stably (panel b).




Fig. 2: Hessian eigenvalue trajectories (left) and training loss (right) for deep linear networks with depths $L = 2$ and $L = 3$.

Sweeping $L = 2-7$ (panel c), the empirical gap $\bar{\lambda}_{\text{dom}}/\bar{\lambda}_{\text{bulk}}$ tracks the $\text{Gap} = L$ prediction: $\bar{\lambda}_{\text{dom}}$ grows linearly in L while $\bar{\lambda}_{\text{bulk}}$ remains nearly constant. This suggests that the depth-dependent bifurcation of Theorem III.5 is not an artifact of linearity but reflects a more general structural property of deep networks.




Fig. 3: (a) Eigenvalues evolution ($L = 3$, tanh) showing spectral bifurcation. (b) Training loss curve. (c) Spectral gap $\bar{\lambda}_{\text{dom}}/\bar{\lambda}_{\text{bulk}}$ vs. depth $L = 2-7$, where $\bar{\lambda}_{\text{dom}}$ and $\bar{\lambda}_{\text{bulk}}$ are the mean eigenvalues of the dominant and bulk spaces, respectively; empirical ratio (red) matches $\text{Gap} = L$ (dashed), $\bar{\lambda}_{\text{dom}}$ (blue) grows linearly, $\bar{\lambda}_{\text{bulk}}$ (green) stays nearly constant.

V. CONCLUSION

In this work, we conduct a theoretical analysis of the spectral bifurcation phenomenon in the Hessian of deep linear neural networks, and demonstrate that this phenomenon is an intrinsic property of the model architecture rather than a consequence of data imbalance. Our main result (Theorem III.5) shows that the Hessian exhibits a two-cluster eigenvalue structure with a gap that scales linearly with the network depth L , even when the data covariance is perfectly balanced. The key insight is that directions within the effective parameter space accumulate curvature contributions from all L layers, while directions outside this space only contribute through a single layer, leading to an L -fold gap between the dominant and bulk eigenvalues. This depth-induced conditioning represents an important factor contributing to the ill-conditioned optimization landscape in deep neural networks, independent of data imbalance, suggesting that both data characteristics and model depth should be considered when designing optimization strategies.

REFERENCES

- [1] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, “How to escape saddle points efficiently,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 1724–1732.
- [2] Z. Yao, A. Gholami, S. Shen, K. Mustafa, and M. Mahoney, “AdaHessian: An adaptive second order optimizer for machine learning,” in *AAAI Conference on Artificial Intelligence*, 2021.
- [3] H. Liu, Z. Li, D. Hall, P. Liang, and T. Ma, “Sophia: A scalable stochastic second-order optimizer for language model pre-training,” *arXiv preprint arXiv:2305.14342*, 2023.
- [4] Y. Zhang, C. Chen, Z. Li, T. Ding, Z. Wu, Y. Ye, Z. He, L. Sun, and Z. Yu, “Adam-mini: Use 1/2 gpu memory of adam with better performance,” *arXiv preprint arXiv:2406.16793*, 2024.
- [5] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, “On large-batch training for deep learning: Generalization gap and sharp minima,” in *International Conference on Learning Representations*, 2017.
- [6] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, “Visualizing the loss landscape of neural nets,” in *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [7] T. Garipov, P. Izmailov, D. Podoprikin, D. P. Vetrov, and A. G. Wilson, “Loss surfaces, mode connectivity, and fast ensembling of dnns,” in *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [8] L. Sagun, U. Evci, V. U. Guney, Y. Dauphin, and L. Bottou, “Empirical analysis of the hessian of over-parametrized neural networks,” *arXiv preprint arXiv:1706.04454*, 2017.
- [9] G. Gur-Ari, D. A. Roberts, and E. Dyer, “Gradient descent happens in a tiny subspace,” *arXiv preprint arXiv:1812.04754*, 2018.
- [10] B. Ghorbani, S. Krishnan, and Y. Xiao, “An investigation into neural net optimization via hessian eigenvalue density,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 2232–2241.
- [11] Z. Yao, A. Gholami, K. Keutzer, and M. W. Mahoney, “Pyhessian: Neural networks through the lens of the hessian,” in *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 2020, pp. 581–590.
- [12] J. Cohen, S. Kaur, Y. Li, J. Z. Kolter, and A. Talwalkar, “Gradient descent on neural networks typically occurs at the edge of stability,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [13] M. Song, K. Ahn, and C. Yun, “Does sgd really happen in tiny subspaces?” 2024.
- [14] C. H. Martin and M. W. Mahoney, “Implicit self-regularization in deep neural networks: Evidence from random matrix theory and implications for learning,” *Journal of Machine Learning Research (JMLR)*, vol. 22, no. 165, pp. 1–73, 2021.
- [15] L. Hodgkinson, Z. Wang, and M. W. Mahoney, “Models of heavy-tailed mechanistic universality,” *arXiv preprint arXiv:2506.03470*, 2025.
- [16] J. Pennington and Y. Bahri, “Geometry of neural network loss surfaces via random matrix theory,” in *International Conference on Machine Learning (ICML)*, 2017, pp. 2798–2806.
- [17] C. Louart, Z. Liao, and R. Couillet, “A random matrix approach to neural networks,” *The Annals of Applied Probability*, vol. 28, no. 2, pp. 1190–1248, 2018.
- [18] J. Pennington and P. Worah, “The spectrum of the fisher information matrix of a single-hidden-layer neural network,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [19] V. Pappas, “Traces of class/cross-class structure in deep network spectra,” *Journal of Machine Learning Research*, vol. 21, no. 252, pp. 1–64, 2020.
- [20] S. Arora, N. Cohen, N. Golowich, and W. Hu, “A convergence analysis of gradient descent for deep linear neural networks,” in *International Conference on Learning Representations*, 2019.
- [21] S. P. Singh, G. Bachmann, and T. Hofmann, “Analytic insights into structure and rank of neural network hessian maps,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 23 914–23 927, 2021.
- [22] B. Bordelon and C. Pehlevan, “Deep linear network training dynamics from random initialization: Data, width, depth, and hyperparameter transfer,” in *International Conference on Machine Learning*, 2025.
- [23] J. Zhao, S. P. Singh, and A. Lucchi, “Theoretical characterisation of the gauss-newton conditioning in neural networks,” in *NeurIPS*, 2024.
- [24] S. Deng, B. Liao, Z. Ouyang, T. Pang, and Y. Yang, “Depth, not data: An analysis of hessian spectral bifurcation,” available on arxiv, 2026.

APPENDIX A
NOTATIONS

In this section, we first recall our notation and foundational definitions that will be used throughout the theoretical analysis.

We follow the standard linear algebra conventions. That is, scalars are denoted by lowercase letters (e.g., λ, η), vectors by bold lowercase letters (e.g., \mathbf{x}, \mathbf{y}), and matrices by capital letters (e.g., W, H). Let I_k denote the identity matrix with rank k . Moreover, we introduce the zero-padding relation \simeq . Specifically, for the matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{p \times q}$, if $m \geq p$ and $n \geq q$, we define the padding operator \simeq as follows:

$$A \simeq B \quad \Leftrightarrow \quad A = \begin{bmatrix} B & \mathbf{0}_{p \times (n-q)} \\ \mathbf{0}_{(m-p) \times q} & \mathbf{0}_{(m-p) \times (n-q)} \end{bmatrix},$$

where $\mathbf{0}_{i \times j}$ denotes an $i \times j$ zero matrix.

For matrix calculus, we use the Kronecker product, denoted by \otimes and adopt a row-wise vectorization convention $\text{vec}_r(\cdot)$.

Our analysis is based on a Deep Linear Network of depth L , parameterized by the set of weight matrices $W := \{W^1, \dots, W^L\}$. Let d_l denote the width of the l -th layer, where d_0 and d_L correspond to the input and output dimensions, respectively. We further define $d_* = \min\{d_0, d_L\}$.

To represent the Deep Linear Network, we use the notation $W^{k:l}$. For $k > l$, this denotes the forward product $W^k W^{k-1} \dots W^l$, while for $k < l$, it denotes the backward transpose product $(W^k)^\top \dots (W^l)^\top$. Therefore, the Deep Linear Network can be expressed as $F_W(\mathbf{x}) = W^{L:1} \mathbf{x}$.

We characterize the data distribution \mathcal{D} over the $d_L \times d_0$ matrices. The input covariance matrix is denoted by $\Sigma_{xx} = \mathbb{E}[\mathbf{x}\mathbf{x}^\top]$, and the input-output cross-covariance is given by $\Sigma_{yx} = \mathbb{E}[\mathbf{y}\mathbf{x}^\top]$. The effective rank of the target mapping is denoted by r , and we define \mathcal{I}_r as $\mathcal{I}_r \in \mathbb{R}^{d_* \times d_*} \simeq I_r$.

The optimization landscape is characterized by the population squared loss, denoted by $L(W)$. Our study focuses mainly on the Hessian matrix of this loss, denoted as $H_L = \nabla^2 L(W)$. Following the Gauss-Newton decomposition, we decompose Hessian into two parts, the outer-product hessian H_o and the functional hessian H_f . η and t represent the learning rate and the training time step index, respectively.

To formally characterize the spectral bifurcation phenomenon, we evaluate the dynamic progress of the singular values of the weight matrices. Let U and V denote the fixed left and right singular basis vectors determined at initialization, respectively. Furthermore, let $\Sigma_t^{1/L}$ be the dynamic diagonal matrix with entries $\lambda_{i,t}$, which are characterized by a mean magnitude m_t and fluctuation δ_t . To study the spectral bifurcation phenomenon precisely, we partition the eigenvalues of the hessian into distinct clusters based on magnitude: dominant eigenvalues (λ_{dom}), and the bulk eigenvalues (λ_{bulk}). The gap between these two clusters drives our theoretical findings for the network depth, which is shown as following.

APPENDIX B
THEORETICAL RESULTS

We are here to give the proof of our theorem. First we need a few lemma.

Lemma B.1. *Under Assumptions III.1, III.2, and III.3, and assuming $r \leq d_* = \min\{d_0, d_L\}$, the outer-product Hessian has two distinct nonzero eigenvalues $L\mu^{2(L-1)}$ and $\mu^{2(L-1)}$, where the former is exactly L times the latter. Moreover, they occur with multiplicities r^2 and $(d_0 + d_L - 2r)r$, respectively.*

Proof. According to [21], for a deep linear network,

$$H_o = A_o B_o A_o^\top \tag{11}$$

where $B_o = I_{d_L} \otimes \Sigma_{xx} \in \mathbb{R}^{d_L d_0 \times d_L d_0}$, and

$$A_o^\top = (W^{L:2} \otimes I_{d_0}, \dots, W^{L:l+1} \otimes W^{1:l-1}, \dots, I_{d_L} \otimes W^{1:L-1}) \tag{12}$$

Under Assumption III.1, $\Sigma_{xx} \simeq I_{d_*}$. When $d_* = d_0 \leq d_L$, we have $\Sigma_{xx} = I_{d_0}$, so $B_o = I_{d_L} \otimes I_{d_0} = I_{d_L d_0}$. In this case, the eigenvalues of $H_o = A_o A_o^\top$ are identical to those of $A_o^\top A_o$.

Under balanced initialization (Initialization II.1) with Assumption III.3, we have $\Sigma_0^{1/L} \simeq \mu I_r$. The weight matrices take the form:

$$W^L \simeq U \Sigma_0^{1/L}, \quad W^{L-1} \simeq \Sigma_0^{1/L}, \quad \dots, \quad W^2 \simeq \Sigma_0^{1/L}, \quad W^1 \simeq \Sigma_0^{1/L} V^\top \tag{13}$$

where $U \in \mathbb{R}^{d_L \times d_*}$ and $V \in \mathbb{R}^{d_0 \times d_*}$ have orthonormal columns. Since $\Sigma_0^{1/L} \simeq \mu I_r$ has only r nonzero entries, let $U_1 \in \mathbb{R}^{d_L \times r}$ and $V_1 \in \mathbb{R}^{d_0 \times r}$ denote the first r columns of U and V , corresponding to the nonzero singular values.

We first compute the weight matrix products. For $W^{L:2}$:

$$W^{L:2} = W^L W^{L-1} \dots W^2 \simeq (U \Sigma_0^{1/L})(\Sigma_0^{1/L}) \dots (\Sigma_0^{1/L}) = U (\Sigma_0^{1/L})^{L-1} \simeq \mu^{L-1} U_1 \tag{14}$$

Therefore:

$$W^{L:2}(W^{L:2})^\top = \mu^{2(L-1)}U_1U_1^\top \quad (15)$$

Similarly, for $W^{1:L-1}$:

$$W^{1:L-1} = (W^1)^\top(W^2)^\top \dots (W^{L-1})^\top \simeq V((\Sigma_0^{1/L})^\top)^{L-1} \simeq \mu^{L-1}V_1 \quad (16)$$

Therefore:

$$W^{1:L-1}(W^{1:L-1})^\top = \mu^{2(L-1)}V_1V_1^\top \quad (17)$$

For $2 \leq l \leq L-1$:

$$W^{L:l+1} \simeq \mu^{L-l}U_1, \quad W^{1:l-1} \simeq \mu^{l-1}V_1 \quad (18)$$

Therefore:

$$W^{L:l+1}(W^{L:l+1})^\top = \mu^{2(L-l)}U_1U_1^\top, \quad W^{1:l-1}(W^{1:l-1})^\top = \mu^{2(l-1)}V_1V_1^\top \quad (19)$$

Now we compute the Gram matrix $A_o^\top A_o$:

$$\begin{aligned} A_o^\top A_o &= (W^{L:2} \otimes I_{d_0})(W^{L:2} \otimes I_{d_0})^\top + (I_{d_L} \otimes W^{1:L-1})(I_{d_L} \otimes W^{1:L-1})^\top \\ &\quad + \sum_{l=2}^{L-1} (W^{L:l+1} \otimes W^{1:l-1})(W^{L:l+1} \otimes W^{1:l-1})^\top \end{aligned} \quad (20)$$

Using the Kronecker product property $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$:

$$(W^{L:2} \otimes I_{d_0})(W^{L:2} \otimes I_{d_0})^\top = W^{L:2}(W^{L:2})^\top \otimes I_{d_0} = \mu^{2(L-1)}U_1U_1^\top \otimes I_{d_0} \quad (21)$$

$$(I_{d_L} \otimes W^{1:L-1})(I_{d_L} \otimes W^{1:L-1})^\top = I_{d_L} \otimes W^{1:L-1}(W^{1:L-1})^\top = \mu^{2(L-1)}I_{d_L} \otimes V_1V_1^\top \quad (22)$$

$$\begin{aligned} &(W^{L:l+1} \otimes W^{1:l-1})(W^{L:l+1} \otimes W^{1:l-1})^\top \\ &= W^{L:l+1}(W^{L:l+1})^\top \otimes W^{1:l-1}(W^{1:l-1})^\top \\ &= \mu^{2(L-l)}U_1U_1^\top \otimes \mu^{2(l-1)}V_1V_1^\top = \mu^{2(L-1)}U_1U_1^\top \otimes V_1V_1^\top \end{aligned} \quad (23)$$

Substituting back:

$$\begin{aligned} A_o^\top A_o &= \mu^{2(L-1)}U_1U_1^\top \otimes I_{d_0} + \mu^{2(L-1)}I_{d_L} \otimes V_1V_1^\top + (L-2)\mu^{2(L-1)}U_1U_1^\top \otimes V_1V_1^\top \\ &= \mu^{2(L-1)} [U_1U_1^\top \otimes I_{d_0} + I_{d_L} \otimes V_1V_1^\top + (L-2)U_1U_1^\top \otimes V_1V_1^\top] \end{aligned} \quad (24)$$

To analyze the eigenvalues, let $\{u_1, \dots, u_r\}$ be the columns of U_1 , extended to an orthonormal basis $\{u_1, \dots, u_{d_L}\}$ for \mathbb{R}^{d_L} . Similarly, let $\{v_1, \dots, v_r\}$ be the columns of V_1 , extended to an orthonormal basis $\{v_1, \dots, v_{d_0}\}$ for \mathbb{R}^{d_0} . The vectors $u_i \otimes v_j$ for $i = 1, \dots, d_L$ and $j = 1, \dots, d_0$ form an orthonormal basis for $\mathbb{R}^{d_L d_0}$.

Note that $U_1U_1^\top u_i = u_i$ if $i \leq r$ and $U_1U_1^\top u_i = 0$ if $i > r$. Similarly, $V_1V_1^\top v_j = v_j$ if $j \leq r$ and $V_1V_1^\top v_j = 0$ if $j > r$. We compute the action of $A_o^\top A_o$ on these basis vectors for each case:

(Dominant space) If $i \leq r$ and $j \leq r$:

$$\begin{aligned} A_o^\top A_o(u_i \otimes v_j) &= \mu^{2(L-1)} [(U_1U_1^\top u_i) \otimes v_j + u_i \otimes (V_1V_1^\top v_j) + (L-2)(U_1U_1^\top u_i) \otimes (V_1V_1^\top v_j)] \\ &= \mu^{2(L-1)} [u_i \otimes v_j + u_i \otimes v_j + (L-2)u_i \otimes v_j] \\ &= L\mu^{2(L-1)}(u_i \otimes v_j) \end{aligned} \quad (25)$$

This gives eigenvalue $L\mu^{2(L-1)}$ with multiplicity $r \times r = r^2$.

(Bulk space, part 1) If $i \leq r$ and $j > r$:

$$\begin{aligned} A_o^\top A_o(u_i \otimes v_j) &= \mu^{2(L-1)} [(U_1U_1^\top u_i) \otimes v_j + u_i \otimes (V_1V_1^\top v_j) + (L-2)(U_1U_1^\top u_i) \otimes (V_1V_1^\top v_j)] \\ &= \mu^{2(L-1)} [u_i \otimes v_j + u_i \otimes 0 + (L-2)u_i \otimes 0] \\ &= \mu^{2(L-1)}(u_i \otimes v_j) \end{aligned} \quad (26)$$

This gives eigenvalue $\mu^{2(L-1)}$ with multiplicity $r \times (d_0 - r) = r(d_0 - r)$.

(Bulk space, part 2) If $i > r$ and $j \leq r$:

$$\begin{aligned} A_o^\top A_o(u_i \otimes v_j) &= \mu^{2(L-1)} [(U_1 U_1^\top u_i) \otimes v_j + u_i \otimes (V_1 V_1^\top v_j) + (L-2)(U_1 U_1^\top u_i) \otimes (V_1 V_1^\top v_j)] \\ &= \mu^{2(L-1)} [0 \otimes v_j + u_i \otimes v_j + (L-2) \cdot 0 \otimes v_j] \\ &= \mu^{2(L-1)} (u_i \otimes v_j) \end{aligned} \quad (27)$$

This gives eigenvalue $\mu^{2(L-1)}$ with multiplicity $(d_L - r) \times r = (d_L - r)r$.

(Zero space) If $i > r$ and $j > r$:

$$A_o^\top A_o(u_i \otimes v_j) = \mu^{2(L-1)} [0 \otimes v_j + u_i \otimes 0 + (L-2) \cdot 0 \otimes 0] = 0 \quad (28)$$

This gives eigenvalue 0 with multiplicity $(d_L - r) \times (d_0 - r)$.

Combining the two parts of the bulk space, the total multiplicity is:

$$r(d_0 - r) + (d_L - r)r = (d_0 + d_L - 2r)r \quad (29)$$

We conclude that the outer-product Hessian has two distinct nonzero eigenvalues $L\mu^{2(L-1)}$ and $\mu^{2(L-1)}$, occurring with multiplicities r^2 and $(d_0 + d_L - 2r)r$, respectively. The total number of nonzero eigenvalues is $r^2 + (d_0 + d_L - 2r)r = (d_0 + d_L - r)r$. \square

Lemma B.2. Under Assumption III.1 and III.2, if $d_L > d_*$ (i.e., $d_L > d_0$), besides W^L , the eigenvectors of W^k are invariant. If $d_L < d_*$ (i.e., $d_L < d_0$), besides W^1 , the eigenvectors of W^k are invariant. Particularly, if $d_L = d_* = d_0$, then for all W^k , the eigenvectors of W^k are invariant.

Proof. Under Assumption III.1, since $d_* \leq \min\{d_{L-1}, d_{L-2}, \dots, d_1\}$, according to the balanced initialization procedure:

$$\begin{aligned} \forall 1 < k \leq L-1, k \in \mathbb{N}, \quad W_0^k &\simeq \Sigma^{1/L} \\ W_0^1 &\simeq \Sigma^{1/L} V^\top \\ W_0^L &\simeq U \Sigma^{1/L} \end{aligned}$$

Then we consider three cases.

1) For $1 < k < L$, considering the dynamic behavior between step $t = 0$ and step $t = 1$:

$$\begin{aligned} \nabla_{W^k} L(W_0^k) &= W_0^{k+1:L} (W_0^{L:1} \Sigma_{xx} - \Sigma_{yx}) W_0^{1:k-1} \\ &\simeq \Sigma^{(L-k)/L} U^\top (U \Sigma V^\top \Sigma_{xx} - \Sigma_{yx}) V \Sigma^{(k-1)/L} \end{aligned} \quad (30)$$

Under Assumption III.1, $\Sigma_{xx} \simeq I_r$. Since the balanced initialization chooses V such that its column space lies within the support of Σ_{xx} , we have $V^\top \Sigma_{xx} V = I_{d_*}$. Therefore,

$$V^\top \Sigma_{xx} = V^\top \quad (31)$$

which implies $U \Sigma V^\top \Sigma_{xx} = U \Sigma V^\top$.

According to Assumption III.2, $\Sigma_{yx} = UV^\top$. Note that $U^\top U = I_{d_*}$ and $V^\top V = I_{d_*}$. Substituting into Equation (30):

$$\begin{aligned} \nabla_{W^k} L(W_0^k) &\simeq \Sigma^{(L-k)/L} U^\top (U \Sigma V^\top - UV^\top) V \Sigma^{(k-1)/L} \\ &= \Sigma^{(L-k)/L} U^\top U (\Sigma - I_{d_*}) V^\top V \Sigma^{(k-1)/L} \\ &= \Sigma^{(L-k)/L} (\Sigma - I_{d_*}) \Sigma^{(k-1)/L} \\ &= \Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L} \end{aligned}$$

Since Σ is a diagonal matrix, $\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}$ is also diagonal. The update rule $W_{t+1}^k = W_t^k - \eta \nabla_{W^k} L(W_t^k)$ only changes the diagonal elements.

For example, suppose $W_0^k \simeq \Sigma^{1/L} = \begin{pmatrix} \lambda_{1,0} & & \\ & \ddots & \\ & & \lambda_{d_*,0} \end{pmatrix}$, then

$$W_1^k = \begin{pmatrix} \lambda_{1,0} - \eta \lambda_{1,0}^{2L-1} + \eta \lambda_{1,0}^{L-1} & & \\ & \ddots & \\ & & \lambda_{d_*,0} - \eta \lambda_{d_*,0}^{2L-1} + \eta \lambda_{d_*,0}^{L-1} \end{pmatrix}$$

Thus the i -th eigenvalue follows the dynamic

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1} \quad (32)$$

and the eigenvectors remain invariant.

2) For $k = 1$, considering the dynamic behavior between step $t = 0$ and step $t = 1$:

$$\begin{aligned} \nabla_{W^1} L(W_0^1) &= W_0^{2:L} (W_0^{L:1} \Sigma_{xx} - \Sigma_{yx}) \\ &\simeq \Sigma^{(L-1)/L} U^\top (U \Sigma V^\top - UV^\top) \\ &= \Sigma^{(L-1)/L} U^\top U (\Sigma - I_{d_*}) V^\top \\ &= \Sigma^{(L-1)/L} (\Sigma - I_{d_*}) V^\top \\ &= (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}) V^\top \end{aligned}$$

For the update rule $W_{t+1}^1 = W_t^1 - \eta \nabla_{W^1} L(W_t^1)$, since $W_0^1 \simeq \Sigma^{1/L} V^\top$:

$$\begin{aligned} W_1^1 &= \Sigma^{1/L} V^\top - \eta (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}) V^\top \\ &= \left[\Sigma^{1/L} - \eta (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}) \right] V^\top \end{aligned}$$

We can separate V^\top , thus the eigenvalues also follow Equation (32) and the eigenvectors (determined by V^\top) remain invariant.

3) For $k = L$, considering the dynamic behavior between step $t = 0$ and step $t = 1$:

$$\begin{aligned} \nabla_{W^L} L(W_0^L) &= (W_0^{L:1} \Sigma_{xx} - \Sigma_{yx}) W_0^{1:L-1} \\ &\simeq (U \Sigma V^\top - UV^\top) V \Sigma^{(L-1)/L} \\ &= U (\Sigma - I_{d_*}) V^\top V \Sigma^{(L-1)/L} \\ &= U (\Sigma - I_{d_*}) \Sigma^{(L-1)/L} \\ &= U (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}) \end{aligned}$$

Now we analyze when the eigenvectors of W^L are invariant.

Note that if $d_L = d_* = d_0$, then $U \in \mathbb{R}^{d_L \times d_L}$ is a square orthogonal matrix. Since $W_0^L \simeq U \Sigma^{1/L}$:

$$\begin{aligned} W_1^L &= U \Sigma^{1/L} - \eta U (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}) \\ &= U \left[\Sigma^{1/L} - \eta (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L}) \right] \end{aligned}$$

The left factor U can be separated, so the eigenvalues also follow Equation (32) and the eigenvectors (determined by U) remain invariant.

However, if $d_L > d_*$ (i.e., $d_L > d_0$), then $U \in \mathbb{R}^{d_L \times d_*}$ is not square, and the gradient $U (\Sigma^{(2L-1)/L} - \Sigma^{(L-1)/L})$ cannot be factored to preserve the eigenvector structure of W^L .

Similarly, if $d_L < d_*$ (i.e., $d_L < d_0$), then $V \in \mathbb{R}^{d_0 \times d_*}$ with $d_0 > d_L$, and the gradient for W^1 cannot be factored to preserve the eigenvector structure of W^1 .

Therefore, the layer-wise Hessian eigenvectors are invariant for all layers except W^L when $d_L > d_*$, or except W^1 when $d_L < d_*$. When $d_L = d_* = d_0$, all eigenvectors are invariant. \square

Lemma B.3. *Under Assumption III.1 and III.2, for initialization $\lambda_{i,0} \in (0, M]$, $i = 1, 2, \dots, r$, where $M = \max_{i=1,2,\dots,r} \{1, \lambda_{i,0}\}$ and proper $\eta < \min \left\{ \frac{1}{L}, \frac{2}{(2L-1)M^{2L-2}} \right\}$, the asymptotic behavior of $\lambda_{i,t}$ will converge to 1 for all $i = 1, 2, \dots, r$, i.e., $\lim_{t \rightarrow \infty} \lambda_{i,t} = 1$. The remaining $d_* - r$ eigenvalues $\lambda_{i,t}$ for $i = r+1, \dots, d_*$ remain unchanged during training.*

Proof. According to Lemma B.2, the gradient of the loss with respect to W^k involves Σ_{xx} . Under Assumption III.1, $\Sigma_{xx} \simeq I_r$, which means only the first r directions of the input space contribute to the gradient.

More precisely, for $1 < k < L$, the gradient is:

$$\nabla_{W^k} L(W_t^k) = W_t^{k+1:L} (W_t^{L:1} \Sigma_{xx} - \Sigma_{yx}) W_t^{1:k-1}$$

Since $\Sigma_{xx} \simeq I_r$ and the balanced initialization chooses V such that its first r rows $V_1 \in \mathbb{R}^{r \times d_*}$ satisfy $V_1^\top V_1 = I_r$ (with $V_2 = \mathbf{0}$), the effective gradient only updates the first r eigenvalues of $\Sigma^{1/L}$.

For the first r eigenvalues, let $i = 1, 2, \dots, r$ be fixed and arbitrary. The update rule is:

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1} \quad (33)$$

For the remaining eigenvalues $i = r+1, \dots, d_*$, since they correspond to directions orthogonal to the support of Σ_{xx} , the gradient contribution is zero, and thus:

$$\lambda_{i,t+1} = \lambda_{i,t}, \quad \forall i = r+1, \dots, d_*$$

Now we prove convergence for $i = 1, 2, \dots, r$. First, we show that for all $t \in \mathbb{N}$, $\lambda_{i,t}$ is bounded by $(0, M]$. The initialization satisfies this condition. For induction, assume $\lambda_{i,t} \in (0, M]$. We consider the following cases:

Case 1: $\lambda_{i,t} \in (0, 1]$. In this regime, the update rule Equation (33) implies $\lambda_{i,t+1} \geq \lambda_{i,t} > 0$ since $0 < \lambda_{i,t} \leq 1$.

If $\lambda_{i,t+1} > M$, it implies:

$$\eta > \frac{M - \lambda_{i,t}}{\lambda_{i,t}^{L-1}(1 - \lambda_{i,t}^L)} \geq \frac{1 - \lambda_{i,t}}{\lambda_{i,t}^{L-1}(1 - \lambda_{i,t}^L)}$$

Note that $1 - \lambda_{i,t}^L = (1 - \lambda_{i,t})(1 + \lambda_{i,t} + \lambda_{i,t}^2 + \dots + \lambda_{i,t}^{L-1})$. Using the fact $\lambda_{i,t} \leq 1$, therefore $\eta > \frac{1 - \lambda_{i,t}}{\lambda_{i,t}^{L-1}(1 - \lambda_{i,t}^L)}$ implies:

$$\eta > \frac{1}{\lambda_{i,t}^{L-1}(1 + \lambda_{i,t} + \lambda_{i,t}^2 + \dots + \lambda_{i,t}^{L-1})} \geq \frac{1}{L}$$

Since $\eta < \min \left\{ \frac{1}{L}, \frac{2}{(2L-1)M^{2L-2}} \right\}$, thus $0 < \lambda_{i,t+1} \leq M$ as well.

Case 2: $\lambda_{i,t} \in (1, M]$. In this case, the update rule Equation (33) implies $\lambda_{i,t+1} \leq \lambda_{i,t}$ since $\lambda_{i,t} > 1$ and $L \geq 1$. Using the induction assumption $\lambda_{i,t} \leq M$ yields $\lambda_{i,t+1} \leq M$ as well. Moreover, since $\eta < \min \left\{ \frac{1}{L}, \frac{2}{(2L-1)M^{2L-2}} \right\} \leq \frac{1}{M^{2L-2}}$:

$$\begin{aligned} \lambda_{i,t+1} &= \lambda_{i,t}(1 - \eta \lambda_{i,t}^{2L-2} + \eta \lambda_{i,t}^{L-2}) \\ &> \lambda_{i,t}(1 - \eta \lambda_{i,t}^{2L-2}) \geq \lambda_{i,t}(1 - \eta M^{2L-2}) > 0 \end{aligned}$$

Therefore, in both cases, we have shown $\lambda_{i,t+1} \in (0, M]$. By the principle of induction, it can be concluded that for all $t \in \mathbb{N}$, $\lambda_{i,t}$ is bounded by $(0, M]$.

Next, we show that after finite steps $\lambda_{i,t}$ will lie in a neighborhood of 1, namely $[1 - \epsilon, 1 + \epsilon]$, where $0 < \epsilon < 1$ is a constant. Define

$$f(\lambda_{i,t}) = (\lambda_{i,t} - 1)^2, \quad g(\lambda_{i,t}) = \lambda_{i,t}^{L-1} - \lambda_{i,t}^{2L-1}$$

Consider:

$$\begin{aligned} f(\lambda_{i,t+1}) - f(\lambda_{i,t}) &= (\lambda_{i,t+1} - 1)^2 - (\lambda_{i,t} - 1)^2 \\ &= ((\lambda_{i,t} - 1) + \eta g(\lambda_{i,t}))^2 - (\lambda_{i,t} - 1)^2 \\ &= 2\eta g(\lambda_{i,t})(\lambda_{i,t} - 1) + \eta^2 g^2(\lambda_{i,t}) \\ &= -\eta g^2(\lambda_{i,t}) \left(\frac{-2(\lambda_{i,t} - 1)}{g(\lambda_{i,t})} - \eta \right) \end{aligned}$$

If $\lambda_{i,t} = 1$, then by the update rule $\lambda_{i,t+1} = \lambda_{i,t+2} = \dots = 1$. Therefore, suppose $\lambda_{i,t} \neq 1$. Note that $g(1) = 0$ and g is differentiable. By Lagrange's Mean Value Theorem, there exists a point ξ between $\lambda_{i,t}$ and 1 such that

$$g(\lambda_{i,t}) = g(\lambda_{i,t}) - g(1) = g'(\xi)(\lambda_{i,t} - 1)$$

Note that $g'(x) = (L-1)x^{L-2} - (2L-1)x^{2L-2}$, and

$$\eta < \min \left\{ \frac{1}{L}, \frac{2}{(2L-1)M^{2L-2}} \right\} \leq \frac{2}{\max_{x \in (0, M]} |g'(x)|}$$

Since $\eta > 0$ and $g^2(\lambda_{i,t}) > 0$ (since $\lambda_{i,t} \neq 1$), thus:

$$f(\lambda_{i,t+1}) - f(\lambda_{i,t}) = -\eta g^2(\lambda_{i,t}) \left(\frac{2}{g'(\xi)} - \eta \right) < 0$$

So long as $\lambda_{i,t}$ is outside the neighborhood $(1 - \epsilon, 1 + \epsilon)$, i.e., $|\lambda_{i,t} - 1| > \epsilon$, the term $|g(\lambda_{i,t})|$ is bounded away from zero. Let

$$C_\epsilon = \min_{x \in (0, M], |x-1| \geq \epsilon} g(x)^2 > 0$$

Then,

$$f(\lambda_{i,t+1}) - f(\lambda_{i,t}) \leq -\eta C_\epsilon \left(\frac{2}{g'(\xi)} - \eta \right) \leq -C_\epsilon \eta^2 < 0$$

where C_ϵ is a strictly positive constant depending on ϵ .

If $\lambda_{i,t}$ never enters the neighborhood $[1-\epsilon, 1+\epsilon]$, the energy $f(\lambda_{i,t})$ would decrease by at least $C_\epsilon \eta^2$ at every step, eventually becoming negative, which contradicts the definition $f(\lambda) \geq 0$. Therefore, the trajectory must enter the neighborhood $[1-\epsilon, 1+\epsilon]$ in finite steps.

Finally, we prove $\lim_{t \rightarrow \infty} \lambda_{i,t} = 1$. Since after finite steps $\lambda_{i,t}$ enters the neighborhood $[1-\epsilon, 1+\epsilon]$ with $0 < \epsilon < 1$, consider the operator $T(\lambda) = \lambda + \eta g(\lambda)$. We have

$$|T'(\lambda)| = |1 + \eta g'(\lambda)| = |1 + \eta((L-1)\lambda^{L-2} - (2L-1)\lambda^{2L-2})| < 1$$

for $\lambda \in [1-\epsilon, 1+\epsilon]$, which implies that T is a contraction mapping on this neighborhood. By the Banach Fixed Point Theorem, $\lambda_{i,t}$ converges to the unique fixed point 1 as $t \rightarrow \infty$.

Since $i \in \{1, 2, \dots, r\}$ is arbitrary, the proof holds for all $i = 1, 2, \dots, r$, which completes the proof. \square

Lemma B.4. *Under Assumption III.1 and III.2, for initialization $\lambda_{i,0} \in (0, M]$, $i = 1, 2, \dots, r$, where $M = \max_{i=1,2,\dots,r} \{1, \lambda_{i,0}\}$. Suppose the step size satisfies $\eta < \min \left\{ \frac{1}{L}, \frac{1}{M^{2L-2}} \right\}$, then the sequence $\{\lambda_{i,t}\}_{t=0}^\infty$ is strictly bounded away from zero. That is, there exists a constant $C_{\min} > 0$ such that $\lambda_{i,t} \geq C_{\min}$ for all $t \geq 0$ and $i = 1, 2, \dots, r$. Furthermore, there exists $T > 0$ such that for all $t \geq T$, $\min_{i=1,\dots,r} (\lambda_{i,t}^{2L-2})$ has a positive lower bound $\alpha > 0$ that is independent of t .*

Proof. Let $i = 1, 2, \dots, r$ be fixed and arbitrary. According to Lemma B.2, the update rule of the eigenvalues is

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1}$$

As shown in Lemma B.3, $0 < \lambda_{i,t} \leq \lambda_{i,t+1}$ if $\lambda_{i,t} \in (0, 1]$ and $\lambda_{i,t+1} \leq \lambda_{i,t} \leq M$ if $\lambda_{i,t} \in [1, M]$. Therefore, to find a lower bound, we only need to analyze the case $\lambda_{i,t} \in [1, M]$. Note that

$$\lambda_{i,t+1} = \lambda_{i,t} (1 - \eta \lambda_{i,t}^{2L-2} + \eta \lambda_{i,t}^{L-2})$$

Let $h(\lambda_{i,t}) = 1 - \eta \lambda_{i,t}^{2L-2} + \eta \lambda_{i,t}^{L-2}$. Since $0 < \eta < \min \left\{ \frac{1}{L}, \frac{1}{M^{2L-2}} \right\} \leq \frac{1}{M^{2L-2}}$, thus

$$h(\lambda_{i,t}) = 1 - \eta \lambda_{i,t}^{2L-2} + \eta \lambda_{i,t}^{L-2} > 1 - \eta \lambda_{i,t}^{2L-2} > 0$$

Since $h(\lambda_{i,t})$ is a continuous function on the compact set $[1, M]$ and $h(\lambda_{i,t}) > 0$ for all $\lambda_{i,t}$, it attains a strictly positive minimum $\delta_i > 0$. Therefore, for $\lambda_{i,t} \in [1, M]$:

$$\lambda_{i,t+1} = \lambda_{i,t} h(\lambda_{i,t}) \geq \delta_i > 0$$

Thus, let $C_i = \min\{\lambda_{i,0}, \delta_i\}$, then for all t , $\lambda_{i,t} \geq C_i$. Let $C_{\min} = \min_{i=1,2,\dots,r} \{C_i\} > 0$, we have for all $i = 1, 2, \dots, r$ and all t , $\lambda_{i,t} \geq C_{\min} > 0$.

Furthermore, by Lemma B.3, for each $i = 1, 2, \dots, r$, there exists $T_i > 0$ such that for all $t \geq T_i$, $\lambda_{i,t} \in [1-\epsilon, 1+\epsilon]$ for some $0 < \epsilon < 1$. Let $T = \max_{i=1,\dots,r} T_i$. Then for all $t \geq T$ and all $i = 1, 2, \dots, r$:

$$\lambda_{i,t} \geq 1 - \epsilon > 0$$

Therefore, for all $t \geq T$:

$$\min_{i=1,\dots,r} (\lambda_{i,t}^{2L-2}) \geq (1 - \epsilon)^{2L-2} := \alpha > 0$$

This lower bound α depends only on ϵ and L , and is independent of t , which completes the proof. \square

Lemma B.5. *Let H_f be a block matrix partitioned into blocks H_{kl} , where $k, l \in \{1, \dots, L\}$. Then the spectral norm of H_f satisfies*

$$\|H_f\|_2 \leq \sqrt{\sum_{k,l} \|H_{kl}\|_2^2}.$$

Proof. For any unit vector \mathbf{x} with $\|\mathbf{x}\|_2 = 1$, we partition \mathbf{x} conformably with the blocks of H_f as $\mathbf{x} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_L^\top]^\top$. The condition $\|\mathbf{x}\|_2 = 1$ implies that $\sum_{l=1}^L \|\mathbf{x}_l\|_2^2 = 1$.

Let $\mathbf{y} = H_f \mathbf{x}$. The k -th block of \mathbf{y} , denoted by \mathbf{y}_k , is given by:

$$\mathbf{y}_k = \sum_{l=1}^L H_{kl} \mathbf{x}_l.$$

By the triangle inequality and the compatibility of the spectral norm, the norm of \mathbf{y}_k is bounded by:

$$\|\mathbf{y}_k\|_2 = \left\| \sum_{l=1}^L H_{kl} \mathbf{x}_l \right\|_2 \leq \sum_{l=1}^L \|H_{kl} \mathbf{x}_l\|_2 \leq \sum_{l=1}^L \|H_{kl}\|_2 \|\mathbf{x}_l\|_2.$$

We compute the squared norm of the entire vector \mathbf{y} :

$$\|H_f \mathbf{x}\|_2^2 = \|\mathbf{y}\|_2^2 = \sum_{k=1}^L \|\mathbf{y}_k\|_2^2 \leq \sum_{k=1}^L \left(\sum_{l=1}^L \|H_{kl}\|_2 \|\mathbf{x}_l\|_2 \right)^2.$$

Applying the Cauchy-Schwarz inequality to the inner sum:

$$\begin{aligned} \left(\sum_{l=1}^L \|H_{kl}\|_2 \|\mathbf{x}_l\|_2 \right)^2 &\leq \left(\sum_{l=1}^L \|H_{kl}\|_2^2 \right) \left(\sum_{l=1}^L \|\mathbf{x}_l\|_2^2 \right) \\ &= \left(\sum_{l=1}^L \|H_{kl}\|_2^2 \right) \cdot 1. \end{aligned}$$

Substituting this back:

$$\|H_f \mathbf{x}\|_2^2 \leq \sum_{k=1}^L \sum_{l=1}^L \|H_{kl}\|_2^2.$$

Since the spectral norm is defined as $\|H_f\|_2 = \sup_{\|\mathbf{x}\|_2=1} \|H_f \mathbf{x}\|_2$, taking the square root of both sides yields the desired upper bound. \square

Lemma B.6. *Under Assumption III.1 and III.2, when $L_W(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2 < \epsilon$, the 2-norm of functional Hessian $\|H_f\|_2$ is bounded by*

$$\|H_f\|_2 \leq \|\Sigma_t^{1/L}\|_2^{L-2} \sqrt{2L(L-1)r} \epsilon^{1/2} = O(\epsilon^{1/2})$$

Proof. According to [21], the functional Hessian H_f can be written as a block matrix, where

$$\begin{aligned} \forall k < l, \quad H_f^{kl} &= W_t^{k+1:l-1} \otimes W_t^{k-1:1} \Omega^\top W_t^{L:l+1} \\ \forall k > l, \quad H_f^{kl} &= W_t^{k+1:L} \Omega W_t^{1:l-1} \otimes W_t^{k-1:l+1} \end{aligned}$$

and the diagonal entries are 0.

By Lemma B.2, the eigenvectors of the weight matrices are invariant during training. Under balanced initialization with $\Sigma_t^{1/L} = \text{diag}(\lambda_{1,t}, \dots, \lambda_{d_*,t})$, we have:

$$\begin{aligned} W_t^k &\simeq \Sigma_t^{1/L}, \quad \text{for } 1 < k < L \\ W_t^1 &\simeq \Sigma_t^{1/L} V^\top \\ W_t^L &\simeq U \Sigma_t^{1/L} \end{aligned}$$

Note that for $k < l$:

$$\begin{aligned} &W_t^{k+1:l-1} \otimes W_t^{k-1:1} \Omega^\top W_t^{L:l+1} \\ &= \Sigma_t^{1/L} \dots \Sigma_t^{1/L} \otimes \Sigma_t^{1/L} \dots \left(\Sigma_t^{1/L} V^\top \right) \Omega^\top \left(U \Sigma_t^{1/L} \right) \dots \Sigma_t^{1/L} \\ &= \Sigma_t^{(l-k-1)/L} \otimes \Sigma_t^{(k-1)/L} V^\top \Omega^\top U \Sigma_t^{(L-l)/L} \end{aligned}$$

And for $k > l$:

$$\begin{aligned} &W_t^{k+1:L} \Omega W_t^{1:l-1} \otimes W_t^{k-1:l+1} \\ &= \left(U \Sigma_t^{1/L} \right) \dots \Sigma_t^{1/L} \Omega \Sigma_t^{1/L} \dots \Sigma_t^{1/L} V^\top \otimes \Sigma_t^{1/L} \dots \Sigma_t^{1/L} \\ &= U \Sigma_t^{(L-k)/L} \Omega V \Sigma_t^{(l-1)/L} \otimes \Sigma_t^{(k-l-1)/L} \end{aligned}$$

For $k < l$, we have $(l-k-1) + (k-1) + (L-l) = L-2$. For $k > l$, we have $(L-k) + (l-1) + (k-l-1) = L-2$. Using the submultiplicativity of the spectral norm $\|AB\|_2 \leq \|A\|_2 \|B\|_2$ and the property $\|A \otimes B\|_2 = \|A\|_2 \|B\|_2$, along with the fact that $\|\Sigma_t^{s/L}\|_2 = \|\Sigma_t^{1/L}\|_2^s$ and $\|U\|_2 = \|V\|_2 = 1$ (since U and V have orthonormal columns), we obtain:

$$\|H_f^{kl}\|_2 \leq \|\Sigma_t^{1/L}\|_2^{L-2} \cdot \|\Omega\|_2$$

Recall that under Assumption III.1, $\Sigma_{xx} \simeq I_r$. Since the balanced initialization chooses V such that its column space lies within the support of Σ_{xx} , we have:

$$\mathbb{E}[\|\mathbf{x}\|_2^2] = \mathbb{E}[\text{Tr}(\mathbf{x}\mathbf{x}^\top)] = \text{Tr}(\mathbb{E}[\mathbf{x}\mathbf{x}^\top]) = \text{Tr}(\Sigma_{xx}) = r$$

Recall that $\Omega = \mathbb{E}[\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}}\mathbf{x}^\top]$, where $\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}} = \hat{\mathbf{y}} - \mathbf{y}$. Using Jensen's inequality, we have:

$$\|\Omega\|_2 = \|\mathbb{E}[\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}}\mathbf{x}^\top]\|_2 \leq \mathbb{E}[\|\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}}\mathbf{x}^\top\|_2] = \mathbb{E}[\|\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}}\|_2\|\mathbf{x}\|_2]$$

Applying the Cauchy-Schwarz inequality for expectations yields:

$$\mathbb{E}[\|\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}}\|_2\|\mathbf{x}\|_2] \leq \sqrt{\mathbb{E}[\|\boldsymbol{\delta}_{\mathbf{x},\mathbf{y}}\|_2^2] \cdot \mathbb{E}[\|\mathbf{x}\|_2^2]} < \sqrt{2r\epsilon}$$

which implies $\|\Omega\|_2 \leq \sqrt{2r\epsilon}$.

Therefore, using $\|H_f\|_2 \leq \sqrt{\sum_{k,l} \|H_f^{kl}\|_2^2}$:

$$\|H_f\|_2 \leq \|\Sigma_t^{1/L}\|_2^{L-2} \sqrt{2L(L-1)r\epsilon} = \|\Sigma_t^{1/L}\|_2^{L-2} \sqrt{2L(L-1)r} \epsilon^{1/2}$$

which completes the proof. \square

Lemma B.7. Under Assumption III.1 and III.2, and assuming the first r rows of V , denoted $V_1 \in \mathbb{R}^{r \times d_*}$, satisfy $V_1 V_1^\top = I_r$, the population loss $L(W_t)$ converges linearly to its minimum possible value. Specifically, let $L_{\min} = \frac{1}{2} \mathbb{E}[\|\mathbf{y}\|_2^2] - \sum_{i=r+1}^{d_*} \lambda_{i,0}^L - \frac{r}{2}$ be the irreducible error, the excess loss $\tilde{L}_t = L(W_t) - L_{\min}$ satisfies:

$$\tilde{L}_t \leq \tilde{L}_0 \cdot e^{-2L\alpha\eta t} = O(e^{-2L\alpha\eta t}) \quad (34)$$

where α is a positive lower bound of $\min_{i=1,\dots,r,t} (\lambda_{i,t}^{2L-2})$.

Proof. Recall that the population loss is $L(W_t) = \frac{1}{2} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\|\mathbf{y} - \hat{\mathbf{y}}_t\|_2^2]$. Thus,

$$\begin{aligned} L(W_t) &= \frac{1}{2} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\mathbf{y}^\top \mathbf{y} - 2\mathbf{y}^\top \hat{\mathbf{y}}_t + \hat{\mathbf{y}}_t^\top \hat{\mathbf{y}}_t] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\|\mathbf{y}\|_2^2] - \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\text{Tr}(\mathbf{y}^\top \hat{\mathbf{y}}_t)] + \frac{1}{2} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\text{Tr}(\hat{\mathbf{y}}_t \hat{\mathbf{y}}_t^\top)] \end{aligned}$$

Note that

$$\begin{aligned} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\text{Tr}(\mathbf{y}^\top \hat{\mathbf{y}}_t)] &= \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\text{Tr}(\mathbf{y}^\top W_t^{L:1} \mathbf{x})] \\ &= \text{Tr}(W_t^{L:1} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\mathbf{x}\mathbf{y}^\top]) \\ &= \text{Tr}(W_t^{L:1} \Sigma_{yx}^\top) \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\text{Tr}(\hat{\mathbf{y}}_t \hat{\mathbf{y}}_t^\top)] &= \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\text{Tr}(W_t^{L:1} \mathbf{x}\mathbf{x}^\top (W_t^{L:1})^\top)] \\ &= \text{Tr}(W_t^{L:1} \mathbb{E}_{\mathbf{x},\mathbf{y} \sim \mu} [\mathbf{x}\mathbf{x}^\top] (W_t^{L:1})^\top) \\ &= \text{Tr}(W_t^{L:1} \Sigma_{xx} (W_t^{L:1})^\top) \end{aligned}$$

By Lemma B.2, the eigenvectors of the weight matrices are invariant during training. Thus, $W_t^{L:1} = U \Sigma_t V^\top$, where $\Sigma_t = \text{diag}(\lambda_{1,t}^L, \dots, \lambda_{d_*,t}^L)$.

Under Assumption III.2, $\Sigma_{yx} = UV^\top$, so $\Sigma_{yx}^\top = VU^\top$. Since $V^\top V = I_{d_*}$ and $U^\top U = I_{d_*}$:

$$\text{Tr}(W_t^{L:1} \Sigma_{yx}^\top) = \text{Tr}(U \Sigma_t V^\top V U^\top) = \text{Tr}(U \Sigma_t U^\top) = \text{Tr}(\Sigma_t) = \sum_{i=1}^{d_*} \lambda_{i,t}^L$$

Under Assumption III.1, $\Sigma_{xx} \simeq I_r$. Partition $V \in \mathbb{R}^{d_0 \times d_*}$ by rows as $V = \begin{pmatrix} V_1 \\ V_2 \end{pmatrix}$ where $V_1 \in \mathbb{R}^{r \times d_*}$. Then:

$$V^\top \Sigma_{xx} V = V_1^\top V_1$$

Under the alignment condition $V_1 V_1^\top = I_r$, the matrix $P := V_1^\top V_1 \in \mathbb{R}^{d_* \times d_*}$ is a projection matrix with rank r . Choosing coordinates such that $P = \text{diag}(I_r, \mathbf{0}_{(d_*-r) \times (d_*-r)})$, we have:

$$\begin{aligned} \text{Tr}(W_t^{L:1} \Sigma_{xx} (W_t^{L:1})^\top) &= \text{Tr}(U \Sigma_t V^\top \Sigma_{xx} V \Sigma_t U^\top) \\ &= \text{Tr}(\Sigma_t P \Sigma_t) = \sum_{i=1}^r \lambda_{i,t}^{2L} \end{aligned}$$

Substituting back into the loss equation:

$$L(W_t) = \frac{1}{2} \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \mu} [\|\mathbf{y}\|^2] - \sum_{i=1}^{d_*} \lambda_{i,t}^L + \frac{1}{2} \sum_{i=1}^r \lambda_{i,t}^{2L}$$

By Lemma B.3, the first r eigenvalues $\lambda_{1,t}, \dots, \lambda_{r,t}$ are updated according to the gradient dynamics and converge to 1, while the remaining $d_* - r$ eigenvalues $\lambda_{r+1,t}, \dots, \lambda_{d_*,t}$ remain at their initial values $\lambda_{r+1,0}, \dots, \lambda_{d_*,0}$.

Separating the terms:

$$L(W_t) = \frac{1}{2} \mathbb{E}[\|\mathbf{y}\|^2] - \sum_{i=1}^r \lambda_{i,t}^L - \sum_{i=r+1}^{d_*} \lambda_{i,0}^L + \frac{1}{2} \sum_{i=1}^r \lambda_{i,t}^{2L}$$

Completing the square for the first r components:

$$\begin{aligned} L(W_t) &= \frac{1}{2} \mathbb{E}[\|\mathbf{y}\|^2] - \sum_{i=r+1}^{d_*} \lambda_{i,0}^L + \frac{1}{2} \sum_{i=1}^r (\lambda_{i,t}^{2L} - 2\lambda_{i,t}^L + 1) - \frac{r}{2} \\ &= \frac{1}{2} \mathbb{E}[\|\mathbf{y}\|^2] - \sum_{i=r+1}^{d_*} \lambda_{i,0}^L - \frac{r}{2} + \frac{1}{2} \sum_{i=1}^r (1 - \lambda_{i,t}^L)^2 \end{aligned}$$

The minimum loss is achieved when $\lambda_{i,t}^L = 1$ for $i = 1, \dots, r$:

$$L_{\min} = \frac{1}{2} \mathbb{E}[\|\mathbf{y}\|^2] - \sum_{i=r+1}^{d_*} \lambda_{i,0}^L - \frac{r}{2}$$

Thus, the excess loss is:

$$\tilde{L}_t = L(W_t) - L_{\min} = \frac{1}{2} \sum_{i=1}^r (1 - \lambda_{i,t}^L)^2$$

According to Lemma B.2, for each $i = 1, \dots, r$, the eigenvalue update rule is:

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1}$$

Note that the function $f(x) = x^L$ is convex when $x > 0$. Thus, by the definition of convex function, for all $1 \leq i \leq r$:

$$\lambda_{i,t+1}^L \geq \lambda_{i,t}^L + L \lambda_{i,t}^{L-1} (\lambda_{i,t+1} - \lambda_{i,t}) = \lambda_{i,t}^L + L \eta \lambda_{i,t}^{L-1} (\lambda_{i,t}^{L-1} - \lambda_{i,t}^{2L-1})$$

which implies:

$$\begin{aligned} 1 - \lambda_{i,t+1}^L &\leq 1 - [\lambda_{i,t}^L + L \eta \lambda_{i,t}^{L-1} (\lambda_{i,t}^{L-1} - \lambda_{i,t}^{2L-1})] \\ &= (1 - \lambda_{i,t}^L) - L \eta \lambda_{i,t}^{2L-2} (1 - \lambda_{i,t}^L) \\ &= (1 - L \eta \lambda_{i,t}^{2L-2}) (1 - \lambda_{i,t}^L) \end{aligned}$$

Therefore,

$$(1 - \lambda_{i,t+1}^L)^2 \leq (1 - L \eta \lambda_{i,t}^{2L-2})^2 (1 - \lambda_{i,t}^L)^2$$

Using the fact $(1 - x)^2 \leq e^{-2x}$ for $x \in [0, 1]$ yields:

$$(1 - \lambda_{i,t+1}^L)^2 \leq e^{-2L \eta \lambda_{i,t}^{2L-2}} (1 - \lambda_{i,t}^L)^2$$

Since α is a positive lower bound of $\min_{i=1, \dots, r, t} (\lambda_{i,t}^{2L-2})$, summing i from 1 to r yields:

$$\tilde{L}_{t+1} = \frac{1}{2} \sum_{i=1}^r (1 - \lambda_{i,t+1}^L)^2 \leq e^{-2L \alpha \eta} \cdot \frac{1}{2} \sum_{i=1}^r (1 - \lambda_{i,t}^L)^2 = e^{-2L \alpha \eta} \tilde{L}_t$$

Repeating this process yields:

$$\tilde{L}_t \leq e^{-2L\alpha\eta t} \tilde{L}_0 = O(e^{-2L\alpha\eta t})$$

which completes the proof. \square

Lemma B.8. *Under Assumptions III.1 and III.2, and assuming $r \leq d_*$, the spectral norm of the functional Hessian $\|H_{f,t}\|_2$ decays exponentially with respect to t . Specifically:*

$$\|H_{f,t}\|_2 = O(\epsilon^{1/2}) = O(e^{-L\alpha\eta t}), \quad (35)$$

where $\alpha = \min_{i=1,\dots,r,t}(\lambda_{i,t}^{2L-2})$ is the minimum of the $2L-2$ power of the eigenvalues over the first r components, which correspond to the effective directions within the support of Σ_{xx} .

Proof. Under Assumption III.1, $\Sigma_{xx} \simeq I_r$, which implies that only the first r eigenvalue directions are effectively updated during gradient descent (as shown in Lemma B.3). The remaining $d_* - r$ eigenvalues remain at their initial values and do not contribute to the loss decrease.

By Lemma B.6, the functional Hessian norm satisfies

$$\|H_{f,t}\|_2 \leq \|\Sigma_t^{1/L}\|_2^{L-2} \sqrt{2L(L-1)r} \epsilon^{1/2}$$

where $\epsilon = L(W_t)$ is the population loss at time t .

By Lemma B.7, the excess loss satisfies

$$\tilde{L}_t \leq \tilde{L}_0 \cdot e^{-2L\alpha\eta t}$$

where $\alpha = \min_{i=1,\dots,r,t}(\lambda_{i,t}^{2L-2})$.

Combining these two results yields

$$\|H_{f,t}\|_2 = O(\epsilon^{1/2}) = O(e^{-L\alpha\eta t})$$

which completes the proof. \square

Lemma B.9. *Under Assumptions III.1 and III.2, and assuming $r \leq d_* = \min\{d_0, d_L\}$, suppose that at time t , $\Sigma_t^{1/L} = \text{diag}(\lambda_{1,t}, \lambda_{2,t}, \dots, \lambda_{r,t}, 0, \dots, 0)$. Suppose $\lambda_{i,t} \in [m_t - \delta_t, m_t + \delta_t]$ for $i = 1, 2, \dots, r$ where $\delta_t > 0$. As long as $\frac{m_t + \delta_t}{m_t - \delta_t} < L^{\frac{1}{2(L-1)}}$, the outer product Hessian $H_{o,t}$ exhibits a three-part spectral structure:*

- 1) **(Dominant space)** r^2 eigenvalues lying in $[L(m_t - \delta_t)^{2(L-1)}, L(m_t + \delta_t)^{2(L-1)}]$.
- 2) **(Bulk space)** $(d_0 + d_L - 2r)r$ eigenvalues lying in $[(m_t - \delta_t)^{2(L-1)}, (m_t + \delta_t)^{2(L-1)}]$.
- 3) **(Zero space)** $(d_L - r)(d_0 - r)$ eigenvalues equal to zero.

The total number of nonzero eigenvalues is $(d_0 + d_L - r)r$.

Proof. According to [21], for a deep linear network at time t ,

$$H_{o,t} = A_{o,t} B_o A_{o,t}^\top$$

where $B_o = I_{d_L} \otimes \Sigma_{xx} \in \mathbb{R}^{d_L d_0 \times d_L d_0}$, and

$$A_{o,t}^\top = (W_t^{L:2} \otimes I_{d_0}, \dots, W_t^{L:l+1} \otimes W_t^{1:l-1}, \dots, I_{d_L} \otimes W_t^{1:L-1})$$

Under Assumption III.1, $\Sigma_{xx} \simeq I_{d_*}$. When $d_* = d_0 \leq d_L$, we have $\Sigma_{xx} = I_{d_0}$, so $B_o = I_{d_L} \otimes I_{d_0} = I_{d_L d_0}$. In this case, the eigenvalues of $H_{o,t} = A_{o,t} B_o A_{o,t}^\top$ are identical to those of $A_{o,t}^\top A_{o,t}$.

By Lemma B.2, the eigenvectors of the weight matrices are invariant during training. Thus, at time t , we have $W_t^{L:1} = U \Sigma_t V^\top$, where $\Sigma_t = \text{diag}(\lambda_{1,t}^L, \dots, \lambda_{r,t}^L, 0, \dots, 0)$ and $\Sigma_t^{1/L} = \text{diag}(\lambda_{1,t}, \dots, \lambda_{r,t}, 0, \dots, 0)$.

Let $U_1 \in \mathbb{R}^{d_L \times r}$ and $V_1 \in \mathbb{R}^{d_0 \times r}$ denote the first r columns of U and V , corresponding to the nonzero singular values. We compute the weight matrix products at time t :

$$\begin{aligned} W_t^{L:2} (W_t^{L:2})^\top &\simeq U_1 \text{diag}(\lambda_{1,t}^{2(L-1)}, \dots, \lambda_{r,t}^{2(L-1)}) U_1^\top \\ W_t^{1:L-1} (W_t^{1:L-1})^\top &\simeq V_1 \text{diag}(\lambda_{1,t}^{2(L-1)}, \dots, \lambda_{r,t}^{2(L-1)}) V_1^\top \end{aligned}$$

And for all $2 \leq l \leq L-1$,

$$\begin{aligned} W_t^{L:l+1} (W_t^{L:l+1})^\top &\simeq U_1 \text{diag}(\lambda_{1,t}^{2(L-l)}, \dots, \lambda_{r,t}^{2(L-l)}) U_1^\top \\ W_t^{1:l-1} (W_t^{1:l-1})^\top &\simeq V_1 \text{diag}(\lambda_{1,t}^{2(l-1)}, \dots, \lambda_{r,t}^{2(l-1)}) V_1^\top \end{aligned}$$

Computing the Gram matrix $A_{o,t}^\top A_{o,t}$:

$$\begin{aligned} A_{o,t}^\top A_{o,t} &= W_t^{L:2} (W_t^{L:2})^\top \otimes I_{d_0} + I_{d_L} \otimes W_t^{1:L-1} (W_t^{1:L-1})^\top \\ &\quad + \sum_{l=2}^{L-1} W_t^{L:l+1} (W_t^{L:l+1})^\top \otimes W_t^{1:l-1} (W_t^{1:l-1})^\top \end{aligned}$$

Let $\Lambda_U^{(l)} = \text{diag}(\lambda_{1,t}^{2(L-l)}, \dots, \lambda_{r,t}^{2(L-l)})$ and $\Lambda_V^{(l)} = \text{diag}(\lambda_{1,t}^{2(l-1)}, \dots, \lambda_{r,t}^{2(l-1)})$. Substituting:

$$A_{o,t}^\top A_{o,t} = U_1 \Lambda_U^{(1)} U_1^\top \otimes I_{d_0} + I_{d_L} \otimes V_1 \Lambda_V^{(L)} V_1^\top + \sum_{l=2}^{L-1} U_1 \Lambda_U^{(l)} U_1^\top \otimes V_1 \Lambda_V^{(l)} V_1^\top$$

To analyze the eigenvalues, let $\{u_1, \dots, u_r\}$ be the columns of U_1 , extended to an orthonormal basis $\{u_1, \dots, u_{d_L}\}$ for \mathbb{R}^{d_L} . Similarly, let $\{v_1, \dots, v_r\}$ be the columns of V_1 , extended to an orthonormal basis $\{v_1, \dots, v_{d_0}\}$ for \mathbb{R}^{d_0} . The vectors $u_i \otimes v_j$ form an orthonormal basis for $\mathbb{R}^{d_L d_0}$.

Note that $U_1 \Lambda_U^{(l)} U_1^\top u_i = \lambda_{i,t}^{2(L-l)} u_i$ if $i \leq r$ and $U_1 \Lambda_U^{(l)} U_1^\top u_i = 0$ if $i > r$. Similarly for $V_1 \Lambda_V^{(l)} V_1^\top$. We compute the action of $A_{o,t}^\top A_{o,t}$ on these basis vectors for each case:

(Dominant space) If $i \leq r$ and $j \leq r$:

$$\begin{aligned} A_{o,t}^\top A_{o,t} (u_i \otimes v_j) &= \lambda_{i,t}^{2(L-1)} (u_i \otimes v_j) + \lambda_{j,t}^{2(L-1)} (u_i \otimes v_j) + \sum_{l=2}^{L-1} \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)} (u_i \otimes v_j) \\ &= \left(\sum_{l=1}^L \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)} \right) (u_i \otimes v_j) \end{aligned}$$

Denote the eigenvalue as $\nu_{i,j} = \sum_{l=1}^L \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)}$. When $i = j$, this equals $L \lambda_{i,t}^{2(L-1)}$. When $i \neq j$ and $\lambda_{i,t} \neq \lambda_{j,t}$, this equals $\frac{\lambda_{i,t}^{2L} - \lambda_{j,t}^{2L}}{\lambda_{i,t}^2 - \lambda_{j,t}^2}$. In all cases, since $\lambda_{i,t}, \lambda_{j,t} \in [m_t - \delta_t, m_t + \delta_t]$:

$$L(m_t - \delta_t)^{2(L-1)} \leq \nu_{i,j} \leq L(m_t + \delta_t)^{2(L-1)}$$

This gives eigenvalue multiplicity $r \times r = r^2$.

(Bulk space, part 1) If $i \leq r$ and $j > r$:

$$A_{o,t}^\top A_{o,t} (u_i \otimes v_j) = \lambda_{i,t}^{2(L-1)} (u_i \otimes v_j) + 0 + 0 = \lambda_{i,t}^{2(L-1)} (u_i \otimes v_j)$$

The eigenvalue is $\nu_{i,j} = \lambda_{i,t}^{2(L-1)} \in [(m_t - \delta_t)^{2(L-1)}, (m_t + \delta_t)^{2(L-1)}]$. This gives eigenvalue multiplicity $r \times (d_0 - r) = r(d_0 - r)$.

(Bulk space, part 2) If $i > r$ and $j \leq r$:

$$A_{o,t}^\top A_{o,t} (u_i \otimes v_j) = 0 + \lambda_{j,t}^{2(L-1)} (u_i \otimes v_j) + 0 = \lambda_{j,t}^{2(L-1)} (u_i \otimes v_j)$$

The eigenvalue is $\nu_{i,j} = \lambda_{j,t}^{2(L-1)} \in [(m_t - \delta_t)^{2(L-1)}, (m_t + \delta_t)^{2(L-1)}]$. This gives eigenvalue multiplicity $(d_L - r) \times r = (d_L - r)r$.

(Zero space) If $i > r$ and $j > r$:

$$A_{o,t}^\top A_{o,t} (u_i \otimes v_j) = 0$$

This gives eigenvalue 0 with multiplicity $(d_L - r) \times (d_0 - r) = (d_L - r)(d_0 - r)$.

Combining the two parts of the bulk space, the total multiplicity is:

$$r(d_0 - r) + (d_L - r)r = (d_0 + d_L - 2r)r$$

The condition $\frac{m_t + \delta_t}{m_t - \delta_t} < L^{\frac{1}{2(L-1)}}$ is equivalent to

$$(m_t + \delta_t)^{2(L-1)} < L(m_t - \delta_t)^{2(L-1)}$$

which guarantees a spectral gap between the dominant space and bulk space:

$$\min_{i,j \leq r} \nu_{i,j} \geq L(m_t - \delta_t)^{2(L-1)} > (m_t + \delta_t)^{2(L-1)} \geq \max_{\substack{i \leq r, j > r \\ \text{or } i > r, j \leq r}} \nu_{i,j}$$

In summary, the outer product Hessian $H_{o,t}$ has:

- r^2 eigenvalues in the dominant space, lying in $[L(m_t - \delta_t)^{2(L-1)}, L(m_t + \delta_t)^{2(L-1)}]$;

- $(d_0 + d_L - 2r)r$ eigenvalues in the bulk space, lying in $[(m_t - \delta_t)^{2(L-1)}, (m_t + \delta_t)^{2(L-1)}]$;
- $(d_L - r)(d_0 - r)$ zero eigenvalues.

The total number of nonzero eigenvalues is $r^2 + (d_0 + d_L - 2r)r = (d_0 + d_L - r)r$. \square

Lemma B.10. *Under Assumptions III.1 and III.2, and assuming $r \leq d_* = \min\{d_0, d_L\}$, suppose that $\{u_1, u_2, \dots, u_{d_L}\}$ and $\{v_1, v_2, \dots, v_{d_0}\}$ are the columns of initialized matrices U (extended to an orthonormal basis of \mathbb{R}^{d_L}) and V (extended to an orthonormal basis of \mathbb{R}^{d_0}), respectively. Then the eigenvectors of the outer product Hessian at time t corresponding to non-zero eigenvalues can be written as*

$$A_{o,t}(u_i \otimes v_j), \quad \text{for } (i \leq r \text{ and } j \leq d_0) \text{ or } (i > r \text{ and } j \leq r)$$

where $A_{o,t}^\top = (W_t^{L:2} \otimes I_{d_0}, \dots, W_t^{L:l+1} \otimes W_t^{1:l-1}, \dots, I_{d_L} \otimes W_t^{1:L-1})$ as defined in [21]. Moreover:

- The r^2 eigenvectors $A_{o,t}(u_i \otimes v_j)$ for $i \leq r$ and $j \leq r$ span the dominant eigenspace.
- The $(d_0 + d_L - 2r)r$ eigenvectors $A_{o,t}(u_i \otimes v_j)$ for $(i \leq r \text{ and } j > r)$ or $(i > r \text{ and } j \leq r)$ span the bulk eigenspace.

Proof. Recall from Lemma B.9 that under Assumption III.1, $\Sigma_{xx} \simeq I_{d_*}$. When $d_* = d_0 \leq d_L$, we have $\Sigma_{xx} = I_{d_0}$, so $B_o = I_{d_L d_0}$. The eigenvalues of $H_{o,t} = A_{o,t} A_{o,t}^\top$ are identical to those of $A_{o,t}^\top A_{o,t}$.

Let $U_1 \in \mathbb{R}^{d_L \times r}$ and $V_1 \in \mathbb{R}^{d_0 \times r}$ denote the first r columns of U and V , corresponding to the nonzero singular values. The Gram matrix $A_{o,t}^\top A_{o,t}$ can be expressed as:

$$A_{o,t}^\top A_{o,t} = U_1 \Lambda_U^{(1)} U_1^\top \otimes I_{d_0} + I_{d_L} \otimes V_1 \Lambda_V^{(L)} V_1^\top + \sum_{l=2}^{L-1} U_1 \Lambda_U^{(l)} U_1^\top \otimes V_1 \Lambda_V^{(l)} V_1^\top$$

where $\Lambda_U^{(l)} = \text{diag}(\lambda_{1,t}^{2(L-l)}, \dots, \lambda_{r,t}^{2(L-l)})$ and $\Lambda_V^{(l)} = \text{diag}(\lambda_{1,t}^{2(l-1)}, \dots, \lambda_{r,t}^{2(l-1)})$.

We show that $u_i \otimes v_j$ are eigenvectors of $A_{o,t}^\top A_{o,t}$. Using the Kronecker product property $(A \otimes B)(x \otimes y) = (Ax) \otimes (By)$, and noting that:

- 1) $U_1 \Lambda_U^{(l)} U_1^\top u_i = \lambda_{i,t}^{2(L-l)} u_i$ if $i \leq r$, and $U_1 \Lambda_U^{(l)} U_1^\top u_i = \mathbf{0}$ if $i > r$.
- 2) $V_1 \Lambda_V^{(l)} V_1^\top v_j = \lambda_{j,t}^{2(l-1)} v_j$ if $j \leq r$, and $V_1 \Lambda_V^{(l)} V_1^\top v_j = \mathbf{0}$ if $j > r$.

We analyze the following cases:

Case 1 (Dominant space): If $i \leq r$ and $j \leq r$, then

$$\begin{aligned} A_{o,t}^\top A_{o,t}(u_i \otimes v_j) &= \lambda_{i,t}^{2(L-1)}(u_i \otimes v_j) + \lambda_{j,t}^{2(L-1)}(u_i \otimes v_j) + \sum_{l=2}^{L-1} \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)}(u_i \otimes v_j) \\ &= \left(\sum_{l=1}^L \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)} \right) (u_i \otimes v_j) \end{aligned}$$

This gives $r \times r = r^2$ eigenvectors spanning the dominant eigenspace.

Case 2 (Bulk space, part 1): If $i \leq r$ and $j > r$, then

$$A_{o,t}^\top A_{o,t}(u_i \otimes v_j) = \lambda_{i,t}^{2(L-1)}(u_i \otimes v_j) + 0 + 0 = \lambda_{i,t}^{2(L-1)}(u_i \otimes v_j)$$

This gives $r \times (d_0 - r) = r(d_0 - r)$ eigenvectors.

Case 3 (Bulk space, part 2): If $i > r$ and $j \leq r$, then

$$A_{o,t}^\top A_{o,t}(u_i \otimes v_j) = 0 + \lambda_{j,t}^{2(L-1)}(u_i \otimes v_j) + 0 = \lambda_{j,t}^{2(L-1)}(u_i \otimes v_j)$$

This gives $(d_L - r) \times r = (d_L - r)r$ eigenvectors.

Case 4 (Zero space): If $i > r$ and $j > r$, then

$$A_{o,t}^\top A_{o,t}(u_i \otimes v_j) = 0$$

This gives $(d_L - r) \times (d_0 - r)$ eigenvectors with zero eigenvalue.

Combining Cases 2 and 3, the bulk eigenspace has $r(d_0 - r) + (d_L - r)r = (d_0 + d_L - 2r)r$ eigenvectors.

By the standard relationship between eigenvectors of $M^\top M$ and MM^\top , the eigenvectors of $H_{o,t} = A_{o,t} A_{o,t}^\top$ corresponding to non-zero eigenvalues are $A_{o,t}(u_i \otimes v_j)$. Among these $(d_0 + d_L - r)r$ eigenvectors:

- r^2 eigenvectors (for $i \leq r$ and $j \leq r$) correspond to the dominant eigenspace with eigenvalues of order $L \lambda_{i,t}^{2(L-1)}$.
- $(d_0 + d_L - 2r)r$ eigenvectors (for $(i \leq r \text{ and } j > r)$ or $(i > r \text{ and } j \leq r)$) correspond to the bulk eigenspace with eigenvalues of order $\lambda_{i,t}^{2(L-1)}$ or $\lambda_{j,t}^{2(L-1)}$.

This completes the proof. \square

Lemma B.11. Under Assumption III.1 and III.2, let $f_{i,j,t}(k)$ be defined as the k th largest value of $\sum_{l=1}^L \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)}$ for $i = 1, \dots, d_L$ and $j = 1, \dots, r$. Then for the true Hessian $H_{L,t}$ at time t , its eigenvalues can be bounded as:

$$f_{i,j,t}(k) - O(e^{-L\alpha\eta t}) \leq \lambda_k(H_{L,t}) \leq f_{i,j,t}(k) + O(e^{-L\alpha\eta t})$$

where λ_k represents the k th largest eigenvalue, and $\alpha = \min_{i=1, \dots, r, t} (\lambda_{i,t}^{2L-2})$ is the minimum of the $2L - 2$ power of the eigenvalues over the first r components, which is a positive constant that is independent of t for sufficiently large t .

Proof. Recall that $H_{L,t} = H_{o,t} + H_{f,t}$. As shown in Lemma B.8, $\|H_{f,t}\|_2 = O(e^{-L\alpha\eta t})$, which implies that $\lambda_{\max}(H_{f,t}) = O(e^{-L\alpha\eta t})$ for the functional Hessian at time t .

Applying Weyl's Inequality yields:

$$\lambda_k(H_{o,t}) - \lambda_{\max}(H_{f,t}) \leq \lambda_k(H_{L,t}) \leq \lambda_k(H_{o,t}) + \lambda_{\max}(H_{f,t})$$

By Lemma B.9, $\lambda_k(H_{o,t}) = f_{i,j,t}(k)$. Therefore:

$$f_{i,j,t}(k) - O(e^{-L\alpha\eta t}) \leq \lambda_k(H_{L,t}) \leq f_{i,j,t}(k) + O(e^{-L\alpha\eta t})$$

which completes the proof. \square

Lemma (Shared Spectral Structure). Under Assumptions III.1 and III.2, all weight matrices W_t^k for $k = 1, \dots, L$ share the same spectral structure $\Sigma_t^{1/L} = \text{diag}(\lambda_{1,t}, \lambda_{2,t}, \dots, \lambda_{d_*,t})$ at any time $t \geq 0$. Specifically:

- 1) For $1 < k < L$: $W_t^k \simeq \Sigma_t^{1/L}$
- 2) For $k = 1$: $W_t^1 \simeq \Sigma_t^{1/L} V^\top$
- 3) For $k = L$: $W_t^L \simeq U \Sigma_t^{1/L}$

where $U \in \mathbb{R}^{d_L \times d_*}$ and $V \in \mathbb{R}^{d_0 \times d_*}$ are the left and right singular vector matrices from the balanced initialization, which remain constant throughout training. Furthermore, the eigenvalues $\lambda_{i,t}$ for $i = 1, \dots, r$ evolve according to:

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1} \quad (36)$$

while the eigenvalues $\lambda_{i,t}$ for $i = r + 1, \dots, d_*$ remain unchanged at their initial values.

Proof. Under Assumption III.1, since $d_* \leq \min\{d_{L-1}, d_{L-2}, \dots, d_1\}$, the balanced initialization procedure yields:

$$\begin{aligned} W_0^k &\simeq \Sigma^{1/L}, \quad \text{for } 1 < k < L \\ W_0^1 &\simeq \Sigma^{1/L} V^\top \\ W_0^L &\simeq U \Sigma^{1/L} \end{aligned}$$

where $A = U \Sigma V^\top$ is the SVD of the initialization matrix A , and $\Sigma^{1/L} = \text{diag}(\lambda_{1,0}, \lambda_{2,0}, \dots, \lambda_{d_*,0})$ with $\lambda_{i,0}$ being the L -th root of the i -th singular value.

By Lemma B.2, the eigenvectors of W_t^k are invariant during training. For $1 < k < L$, the gradient is:

$$\nabla_{W^k} L(W_t^k) = \Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L}$$

which is a diagonal matrix. Therefore, the update $W_{t+1}^k = W_t^k - \eta \nabla_{W^k} L(W_t^k)$ only modifies the diagonal entries, preserving the diagonal structure.

For $k = 1$, the gradient is:

$$\nabla_{W^1} L(W_t^1) = (\Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L}) V^\top$$

Thus:

$$\begin{aligned} W_{t+1}^1 &= W_t^1 - \eta \nabla_{W^1} L(W_t^1) \\ &= \Sigma_t^{1/L} V^\top - \eta (\Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L}) V^\top \\ &= \left[\Sigma_t^{1/L} - \eta (\Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L}) \right] V^\top \\ &= \Sigma_{t+1}^{1/L} V^\top \end{aligned}$$

where the right factor V^\top remains unchanged.

For $k = L$, the gradient is:

$$\nabla_{W^L} L(W_t^L) = U (\Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L})$$

Thus:

$$\begin{aligned}
W_{t+1}^L &= W_t^L - \eta \nabla_{W^L} L(W_t^L) \\
&= U \Sigma_t^{1/L} - \eta U (\Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L}) \\
&= U \left[\Sigma_t^{1/L} - \eta (\Sigma_t^{(2L-1)/L} - \Sigma_t^{(L-1)/L}) \right] \\
&= U \Sigma_{t+1}^{1/L}
\end{aligned}$$

where the left factor U remains unchanged.

In all cases, the updated weight matrix shares the same diagonal structure $\Sigma_{t+1}^{1/L}$, where the i -th diagonal entry evolves according to:

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta \lambda_{i,t}^{2L-1} + \eta \lambda_{i,t}^{L-1}$$

By Lemma B.3, under Assumption III.1, only the first r eigenvalues $\lambda_{1,t}, \dots, \lambda_{r,t}$ are effectively updated since $\Sigma_{xx} \simeq I_r$. The remaining eigenvalues $\lambda_{r+1,t}, \dots, \lambda_{d_*,t}$ correspond to directions orthogonal to the support of Σ_{xx} , and thus their gradients are zero:

$$\lambda_{i,t+1} = \lambda_{i,t} = \lambda_{i,0}, \quad \text{for } i = r+1, \dots, d_*$$

By induction on t , all weight matrices W_t^k share the same spectral structure $\Sigma_t^{1/L}$ at any time $t \geq 0$, which completes the proof. \square

Theorem. III.5(Hessian Bifurcation) Under Assumptions III.1 and III.2, and assuming $r \leq d_* = \min\{d_0, d_L\}$, consider a depth- L deep linear neural network trained with gradient descent with step size $\eta < \min\left\{\frac{1}{L}, \frac{2}{(2L-1)M^{2L-2}}\right\}$. Let $\lambda_{i,t}$ for $i = 1, \dots, r$ denote the effective eigenvalues of the weight matrix $\Sigma_t^{1/L}$ at time t . Suppose that $\lambda_{i,t} \in [m_t - \delta_t, m_t + \delta_t]$ and the condition $(m_t + \delta_t)/(m_t - \delta_t) < L^{\frac{1}{2(L-1)}}$ holds. Then the Hessian $H_{L,t}$ exhibits a three-part spectral structure: leftmargin=*

1) **(Dominant Space)** There exist r^2 eigenvalues of $H_{L,t}$ lying in:

$$\begin{aligned}
&\left[L(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t}), \right. \\
&\quad \left. L(m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t}) \right]
\end{aligned} \tag{37}$$

2) **(Bulk Space)** There exist $(d_0 + d_L - 2r)r$ eigenvalues of $H_{L,t}$ lying in:

$$\begin{aligned}
&\left[(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t}), \right. \\
&\quad \left. (m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t}) \right]
\end{aligned} \tag{38}$$

3) **(Zero Space)** The remaining $(d_L - r)(d_0 - r)$ eigenvalues of $H_{L,t}$ are $O(e^{-L\alpha\eta t})$, converging to zero as $t \rightarrow \infty$.

Moreover, let λ_{dom} denote an arbitrary eigenvalue belonging to the Dominant Space and λ_{bulk} denote an arbitrary eigenvalue belonging to the Bulk Space. Their ratio satisfies:

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} = \Theta(L), \tag{39}$$

where α is a positive constant independent of t , provided t is sufficiently large.

Proof. By Lemma B.9, under the assumption $r \leq d_*$, the outer product Hessian $H_{o,t}$ at time t has $(d_0 + d_L - r)r$ non-zero eigenvalues. These eigenvalues can be indexed by pairs (i, j) and exhibit three different behaviors:

Case 1 (Dominant space): If $i \leq r$ and $j \leq r$, the eigenvalue is

$$\nu_{i,j} = \sum_{l=1}^L \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)}$$

When $i = j$, this equals $L \lambda_{i,t}^{2(L-1)}$. When $i \neq j$ and $\lambda_{i,t} \neq \lambda_{j,t}$, this equals $\frac{\lambda_{i,t}^{2L} - \lambda_{j,t}^{2L}}{\lambda_{i,t}^2 - \lambda_{j,t}^2}$. When $i \neq j$ but $\lambda_{i,t} = \lambda_{j,t}$, this also equals $L \lambda_{i,t}^{2(L-1)}$ by continuity.

Case 2 (Bulk space, part 1): If $i \leq r$ and $j > r$, the eigenvalue is

$$\nu_{i,j} = \lambda_{i,t}^{2(L-1)}$$

Case 3 (Bulk space, part 2): If $i > r$ and $j \leq r$, the eigenvalue is

$$\nu_{i,j} = \lambda_{j,t}^{2(L-1)}$$

Case 4 (Zero space): If $i > r$ and $j > r$, the eigenvalue is zero.

Since $\lambda_{i,t} \in [m_t - \delta_t, m_t + \delta_t]$ for all $i = 1, \dots, r$, we bound the eigenvalues in each case.

For Case 1, since each term $\lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)}$ satisfies

$$(m_t - \delta_t)^{2(L-1)} \leq \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)} \leq (m_t + \delta_t)^{2(L-1)}$$

summing over $l = 1, \dots, L$ yields

$$L(m_t - \delta_t)^{2(L-1)} \leq \nu_{i,j} \leq L(m_t + \delta_t)^{2(L-1)}$$

This gives $r \times r = r^2$ eigenvalues in the dominant space.

For Cases 2 and 3:

$$(m_t - \delta_t)^{2(L-1)} \leq \nu_{i,j} \leq (m_t + \delta_t)^{2(L-1)}$$

Case 2 gives $r \times (d_0 - r) = r(d_0 - r)$ eigenvalues, and Case 3 gives $(d_L - r) \times r = (d_L - r)r$ eigenvalues. Combining these, the bulk space contains

$$r(d_0 - r) + (d_L - r)r = (d_0 + d_L - 2r)r$$

eigenvalues.

For Case 4, there are $(d_L - r) \times (d_0 - r) = (d_L - r)(d_0 - r)$ zero eigenvalues.

The condition $(m_t + \delta_t)/(m_t - \delta_t) < L^{\frac{1}{2(L-1)}}$ is equivalent to

$$(m_t + \delta_t)^{2(L-1)} < L(m_t - \delta_t)^{2(L-1)}$$

which guarantees a spectral gap between the dominant space and the bulk space:

$$\min_{i,j \leq r} \nu_{i,j} \geq L(m_t - \delta_t)^{2(L-1)} > (m_t + \delta_t)^{2(L-1)} \geq \max_{\substack{i \leq r, j > r \\ \text{or } i > r, j \leq r}} \nu_{i,j}$$

By the Gauss-Newton decomposition, $H_{L,t} = H_{o,t} + H_{f,t}$. According to Lemma B.8, the functional Hessian satisfies $\|H_{f,t}\|_2 = O(e^{-L\alpha\eta t})$. Applying Weyl's inequality, for the k th largest eigenvalue:

$$\lambda_k(H_{o,t}) - \|H_{f,t}\|_2 \leq \lambda_k(H_{L,t}) \leq \lambda_k(H_{o,t}) + \|H_{f,t}\|_2$$

Therefore, the eigenvalues of the true Hessian $H_{L,t}$ lie within $O(e^{-L\alpha\eta t})$ of the eigenvalues of the outer product Hessian $H_{o,t}$. Combining with the bounds above:

- The r^2 dominant space eigenvalues of $H_{L,t}$ lie in

$$\left[L(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t}), L(m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t}) \right]$$

- The $(d_0 + d_L - 2r)r$ bulk space eigenvalues of $H_{L,t}$ lie in

$$\left[(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t}), (m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t}) \right]$$

- The remaining $(d_L - r)(d_0 - r)$ eigenvalues are $O(e^{-L\alpha\eta t})$, converging to zero as $t \rightarrow \infty$.

Finally, we establish the ratio between the dominant space eigenvalues and bulk space eigenvalues. For the lower bound:

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} \geq \frac{L(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t})}{(m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t})}$$

For sufficiently large t , the $O(e^{-L\alpha\eta t})$ terms become negligible, yielding

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} \geq \frac{L(m_t - \delta_t)^{2(L-1)}}{(m_t + \delta_t)^{2(L-1)}} = L \left(\frac{m_t - \delta_t}{m_t + \delta_t} \right)^{2(L-1)}$$

For the upper bound:

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} \leq \frac{L(m_t + \delta_t)^{2(L-1)} + O(e^{-L\alpha\eta t})}{(m_t - \delta_t)^{2(L-1)} - O(e^{-L\alpha\eta t})}$$

For sufficiently large t :

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} \leq \frac{L(m_t + \delta_t)^{2(L-1)}}{(m_t - \delta_t)^{2(L-1)}} = L \left(\frac{m_t + \delta_t}{m_t - \delta_t} \right)^{2(L-1)}$$

Since the condition $(m_t + \delta_t)/(m_t - \delta_t) < L^{\frac{1}{2(L-1)}}$ implies that $\left(\frac{m_t + \delta_t}{m_t - \delta_t}\right)^{2(L-1)} < L$, both bounds are $O(L)$. Therefore:

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} = \Theta(L)$$

which completes the proof. \square

Corollary. III.6(Hessian Bifurcation with USI) Under Assumptions III.1, III.2, and III.3, and assuming $r \leq d_* = \min\{d_0, d_L\}$, the Hessian $H_{L,t}$ has exactly two distinct nonzero eigenvalues (up to an exponentially small perturbation). Specifically: leftmargin=*

- 1) **(Dominant Space)** There exist r^2 eigenvalues equal to $L\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})$.
- 2) **(Bulk Space)** There exist $(d_0 + d_L - 2r)r$ eigenvalues equal to $\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})$.
- 3) **(Zero Space)** The remaining $(d_L - r)(d_0 - r)$ eigenvalues are $O(e^{-L\alpha\eta t})$.

Moreover, the ratio between the dominant and bulk eigenvalues is exactly L , i.e.,

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} = L \tag{40}$$

where μ_t denotes the common value of all effective eigenvalues $\lambda_{i,t} = \mu_t$ for $i = 1, \dots, r$ at time t , and α is a positive constant.

Proof. Under Assumption III.3, $\Sigma_0^{1/L} \simeq \mu I_r$, which implies $\lambda_{i,0} = \mu$ for all $i = 1, \dots, r$. By Lemma III.4, the effective eigenvalues evolve according to:

$$\lambda_{i,t+1} = \lambda_{i,t} - \eta\lambda_{i,t}^{2L-1} + \eta\lambda_{i,t}^{L-1}$$

for $i = 1, \dots, r$.

Since all initial eigenvalues are equal, and the update rule is identical for each $i = 1, \dots, r$, we have $\lambda_{1,t} = \lambda_{2,t} = \dots = \lambda_{r,t} =: \mu_t$ for all $t \geq 0$. By Lemma B.3, $\mu_t \rightarrow 1$ as $t \rightarrow \infty$.

By Lemma B.9, the outer product Hessian $H_{o,t}$ has eigenvalues indexed by pairs (i, j) . For $i \leq r$ and $j \leq r$:

$$\nu_{i,j} = \sum_{l=1}^L \lambda_{i,t}^{2(L-l)} \lambda_{j,t}^{2(l-1)} = \sum_{l=1}^L \mu_t^{2(L-l)} \mu_t^{2(l-1)} = \sum_{l=1}^L \mu_t^{2(L-1)} = L\mu_t^{2(L-1)}$$

For $i \leq r$ and $j > r$:

$$\nu_{i,j} = \lambda_{i,t}^{2(L-1)} = \mu_t^{2(L-1)}$$

For $i > r$ and $j \leq r$:

$$\nu_{i,j} = \lambda_{j,t}^{2(L-1)} = \mu_t^{2(L-1)}$$

For $i > r$ and $j > r$:

$$\nu_{i,j} = 0$$

Therefore, the outer product Hessian $H_{o,t}$ has exactly two distinct nonzero eigenvalues:

- $L\mu_t^{2(L-1)}$ with multiplicity r^2 (dominant space)
- $\mu_t^{2(L-1)}$ with multiplicity $r(d_0 - r) + (d_L - r)r = (d_0 + d_L - 2r)r$ (bulk space)

and zero eigenvalues with multiplicity $(d_L - r)(d_0 - r)$ (zero space).

By Lemma B.8, $\|H_{f,t}\|_2 = O(e^{-L\alpha\eta t})$. Applying Weyl's inequality:

$$\lambda_k(H_{o,t}) - O(e^{-L\alpha\eta t}) \leq \lambda_k(H_{L,t}) \leq \lambda_k(H_{o,t}) + O(e^{-L\alpha\eta t})$$

Therefore, the Hessian $H_{L,t}$ has:

- r^2 eigenvalues equal to $L\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})$
- $(d_0 + d_L - 2r)r$ eigenvalues equal to $\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})$
- $(d_L - r)(d_0 - r)$ eigenvalues equal to $O(e^{-L\alpha\eta t})$

The ratio between the dominant and bulk eigenvalues is:

$$\frac{\lambda_{\text{dom}}}{\lambda_{\text{bulk}}} = \frac{L\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})}{\mu_t^{2(L-1)} + O(e^{-L\alpha\eta t})} = L + O(e^{-L\alpha\eta t})$$

For sufficiently large t , this ratio converges to exactly L , which completes the proof. □

APPENDIX C SIMULATIONS

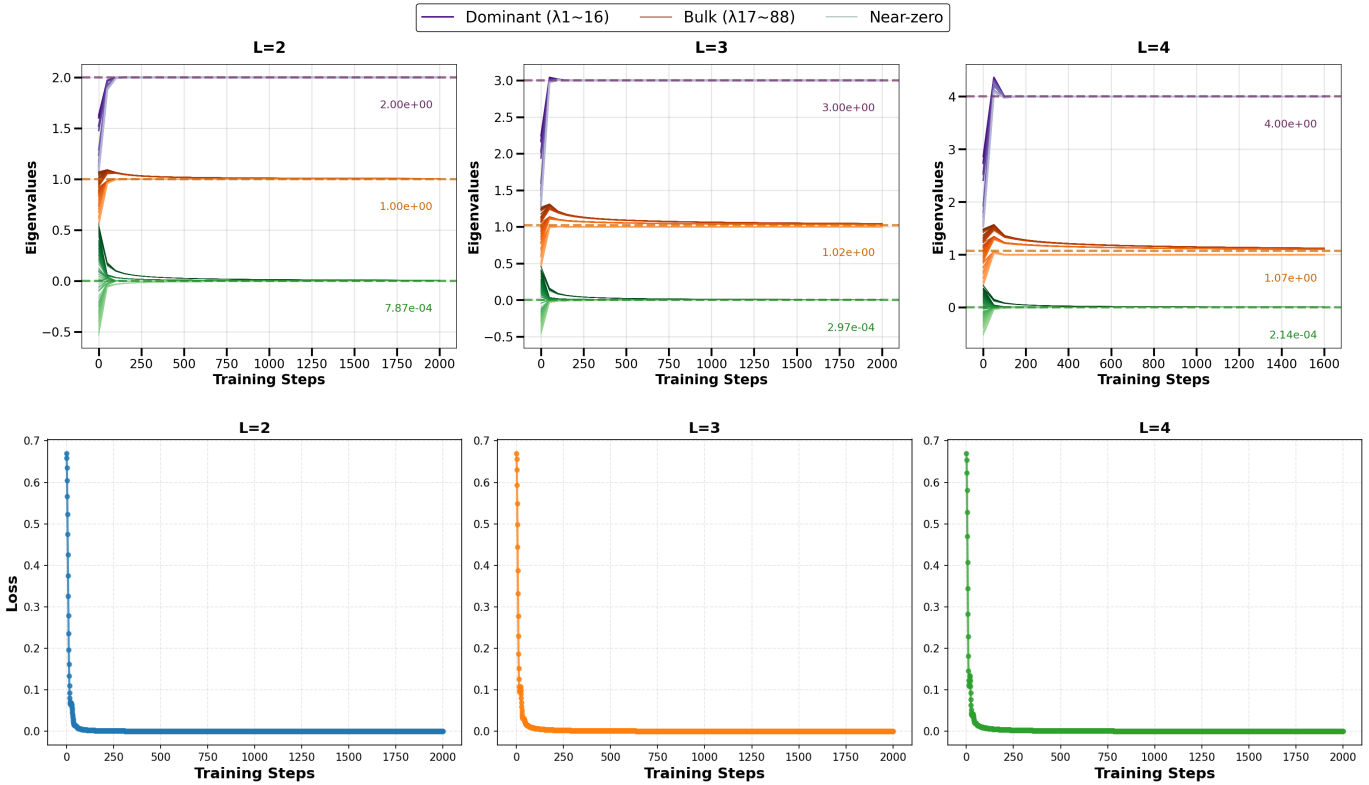


Fig. 4: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

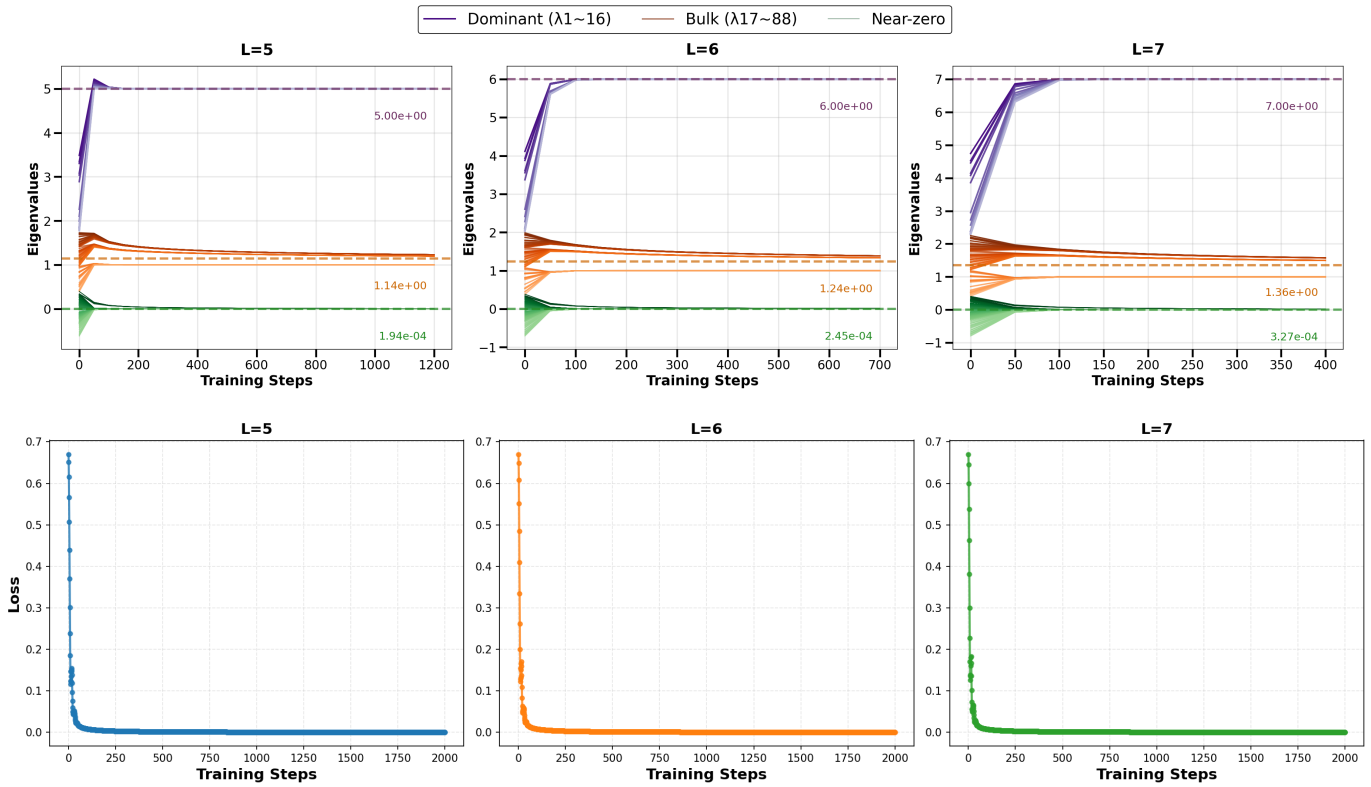


Fig. 5: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

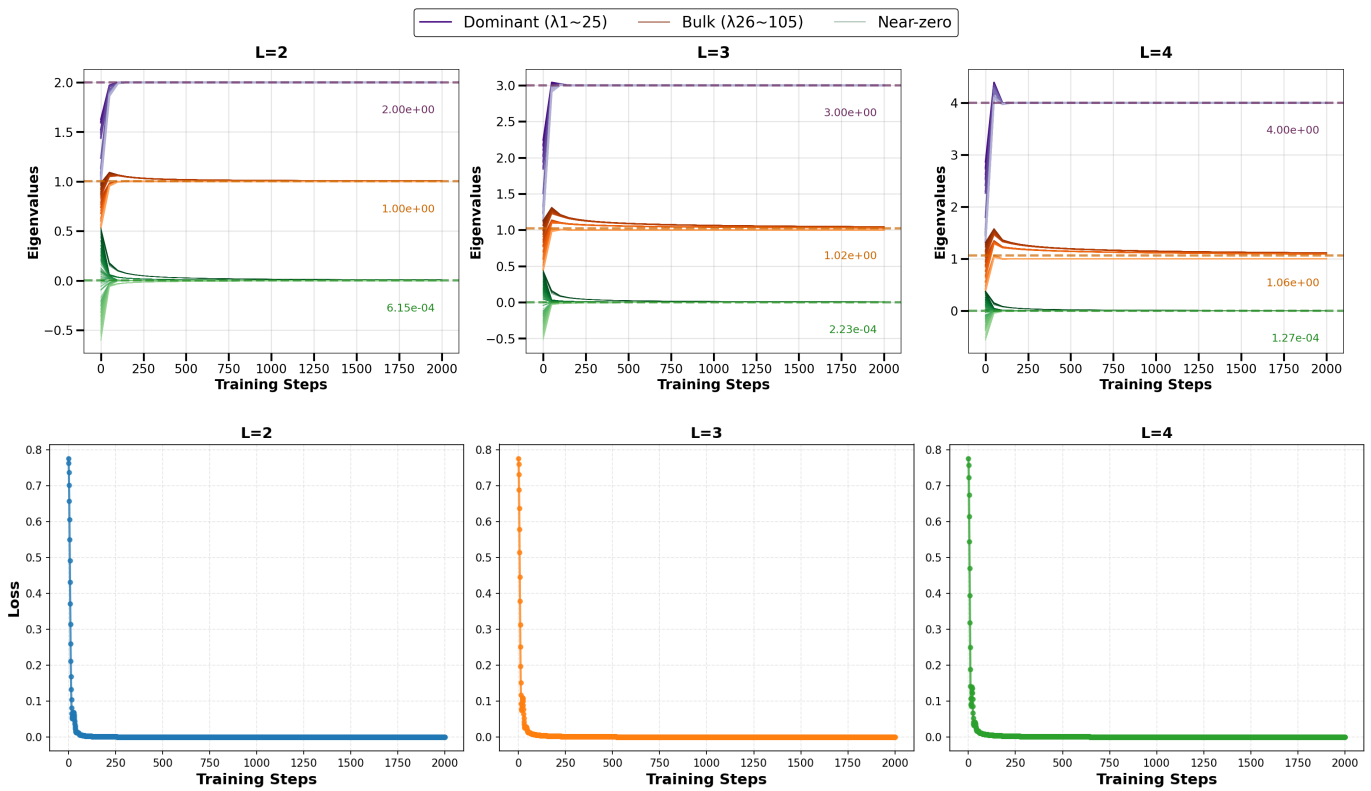


Fig. 6: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

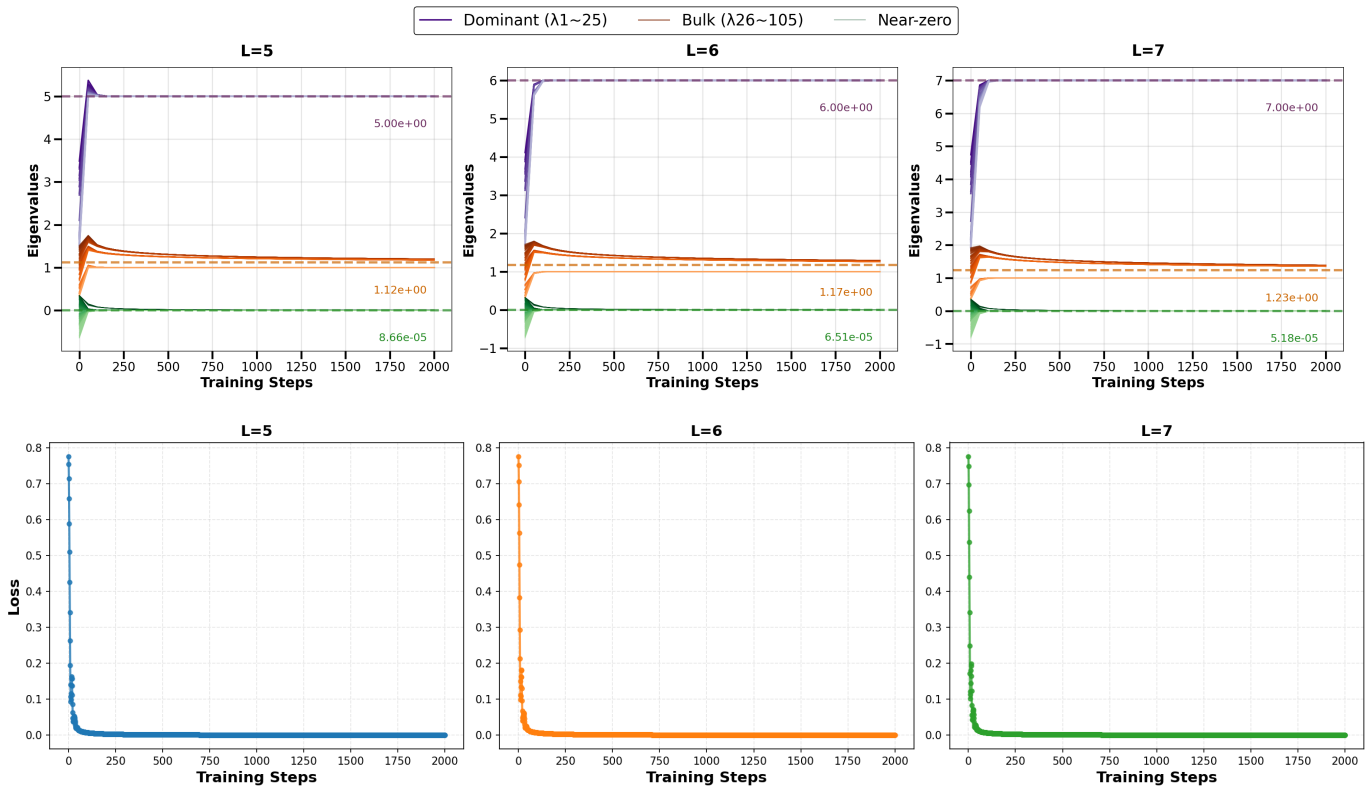


Fig. 7: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

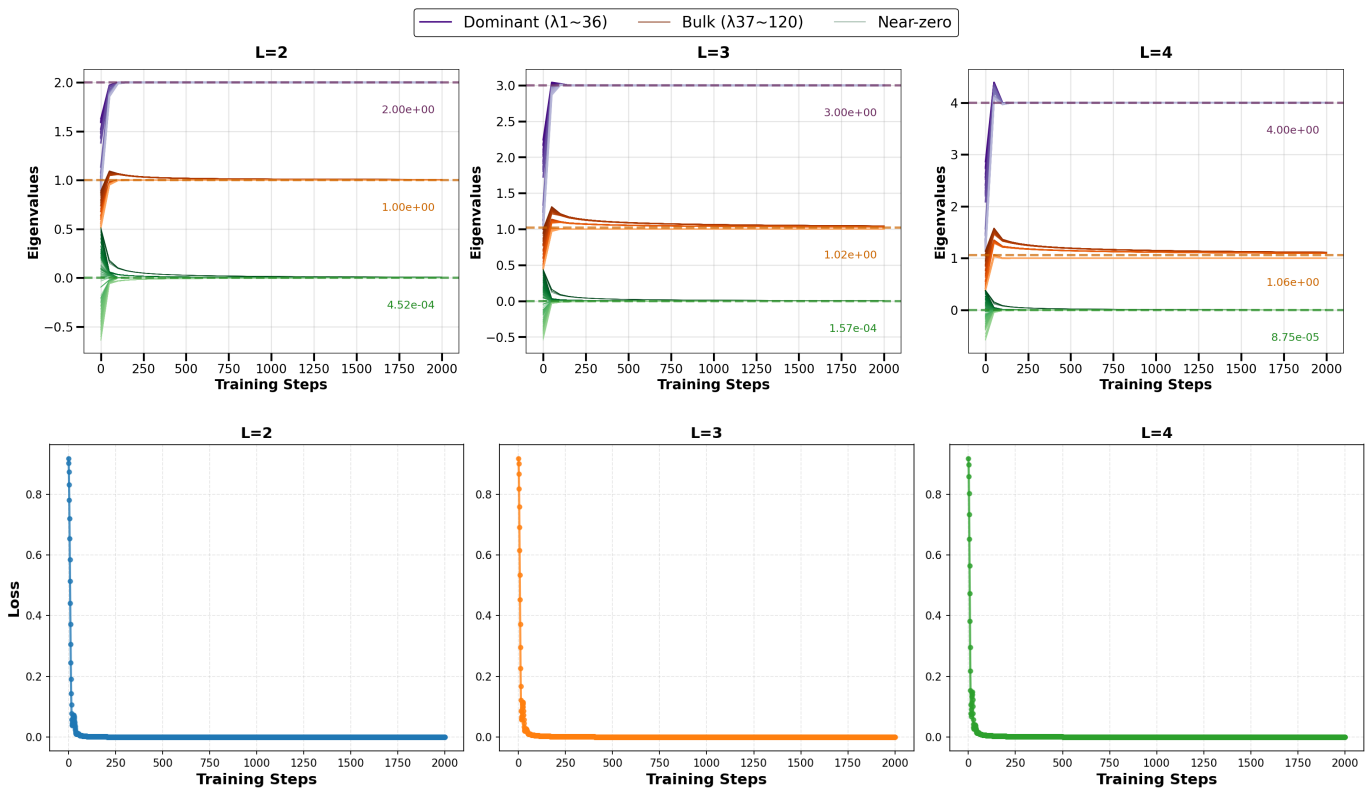


Fig. 8: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

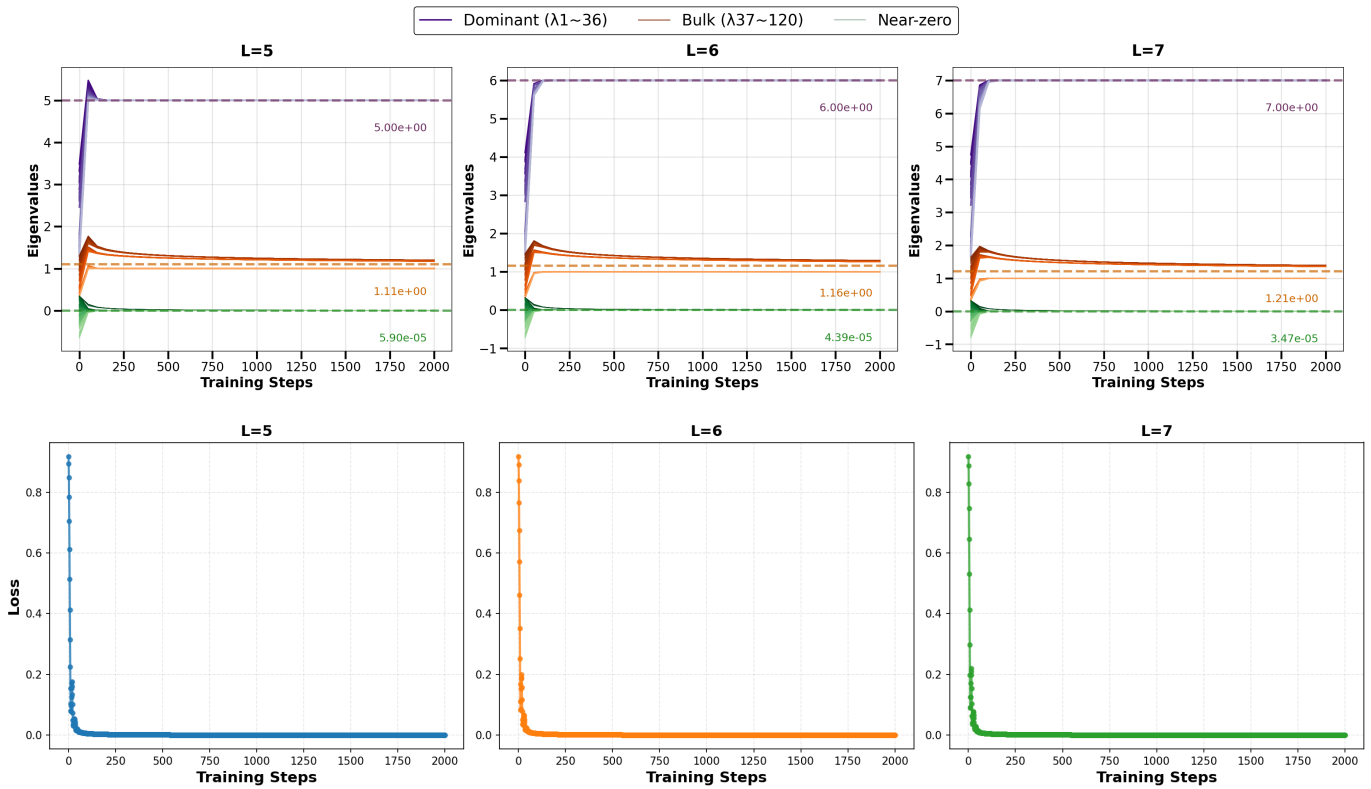


Fig. 9: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

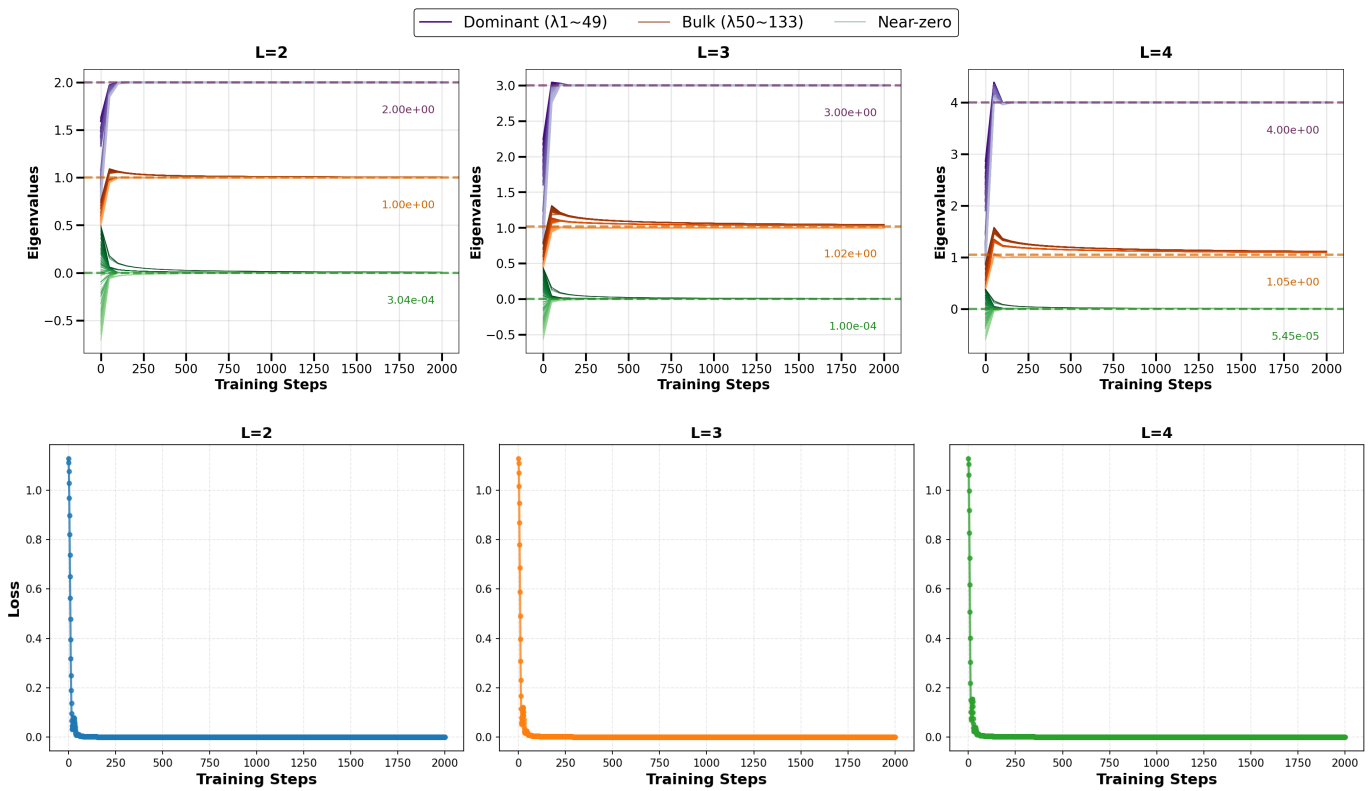


Fig. 10: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

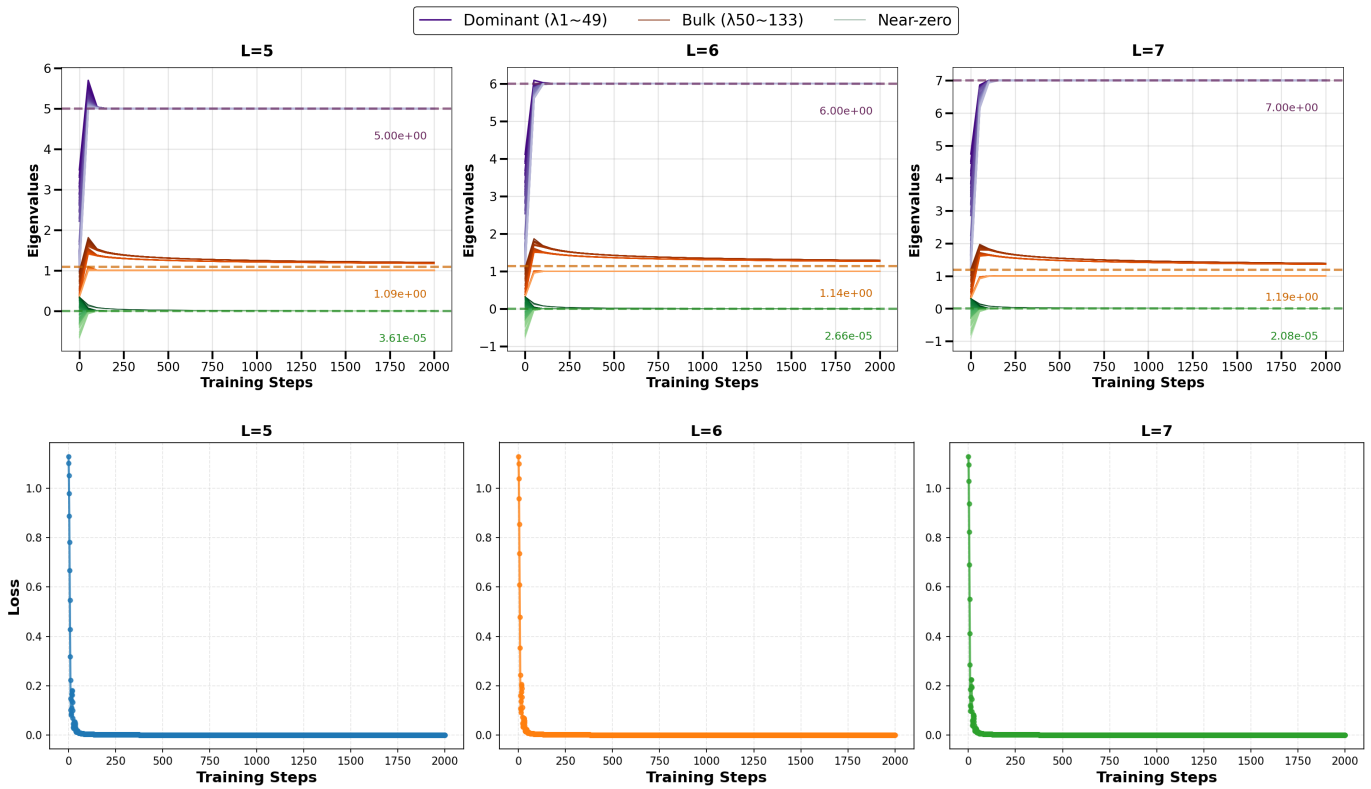


Fig. 11: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

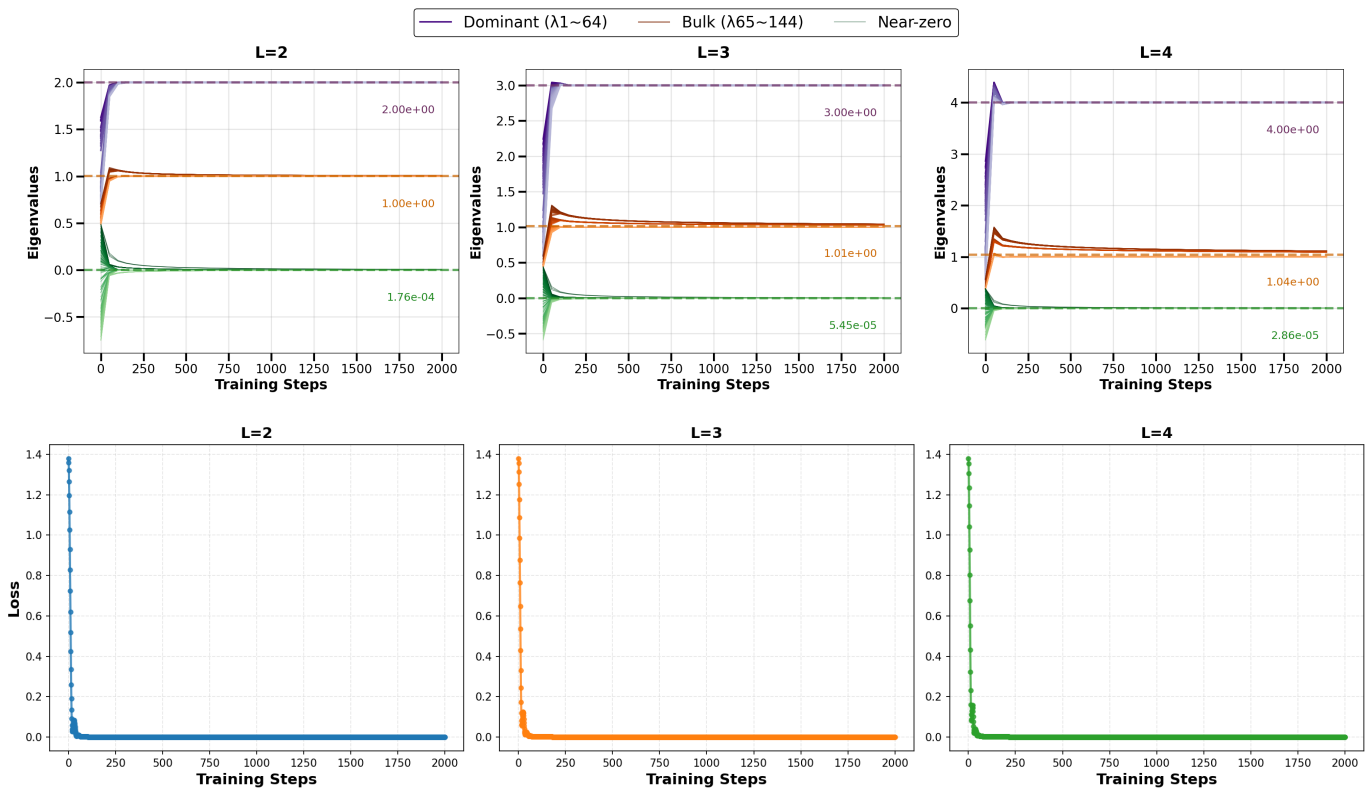


Fig. 12: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

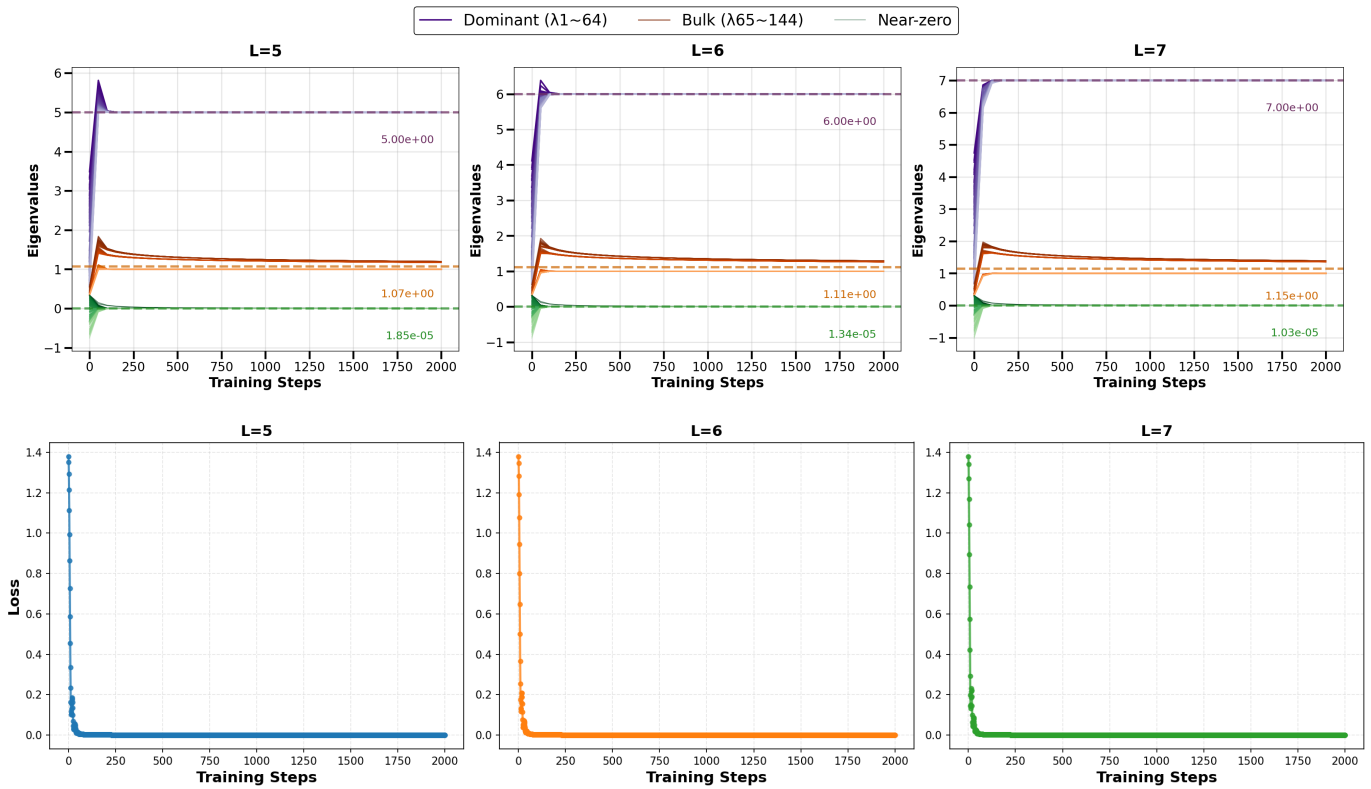


Fig. 13: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

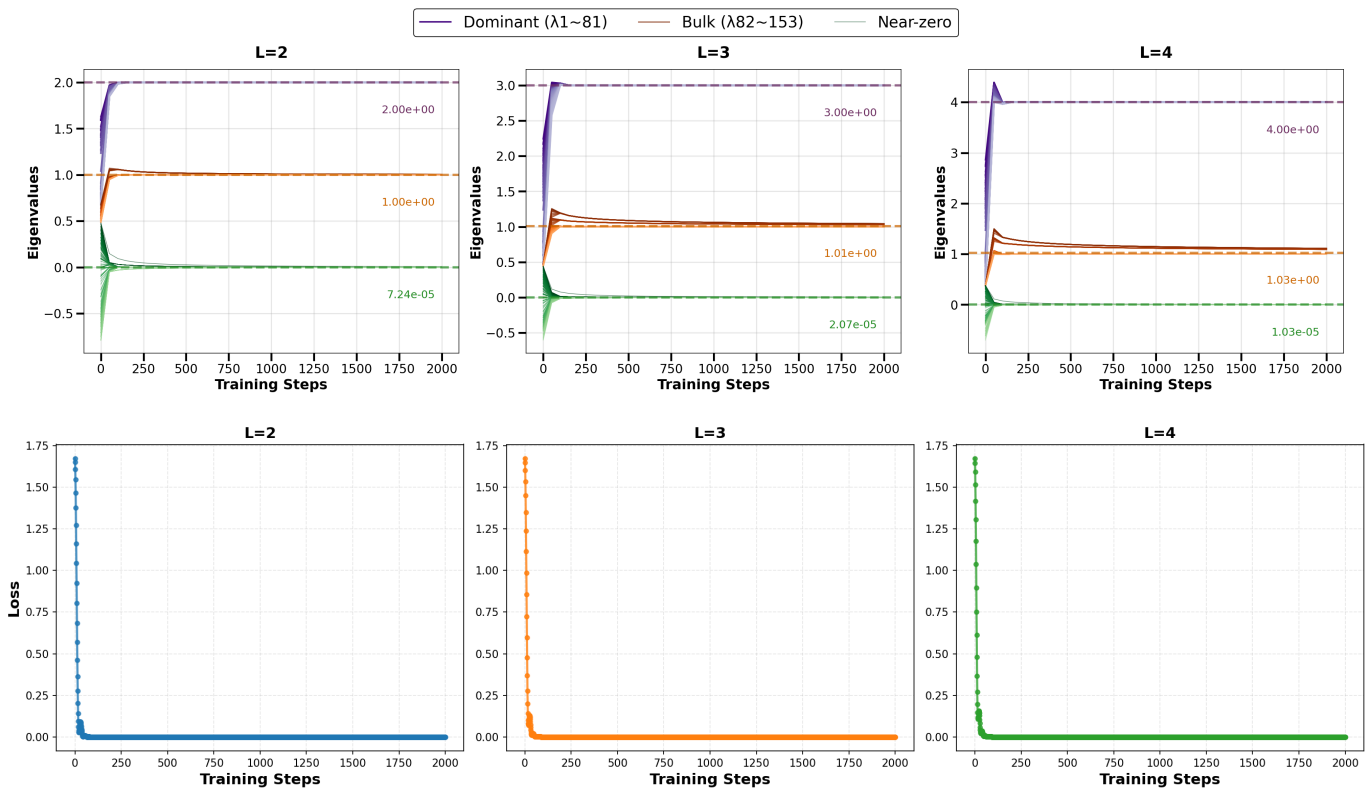


Fig. 14: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

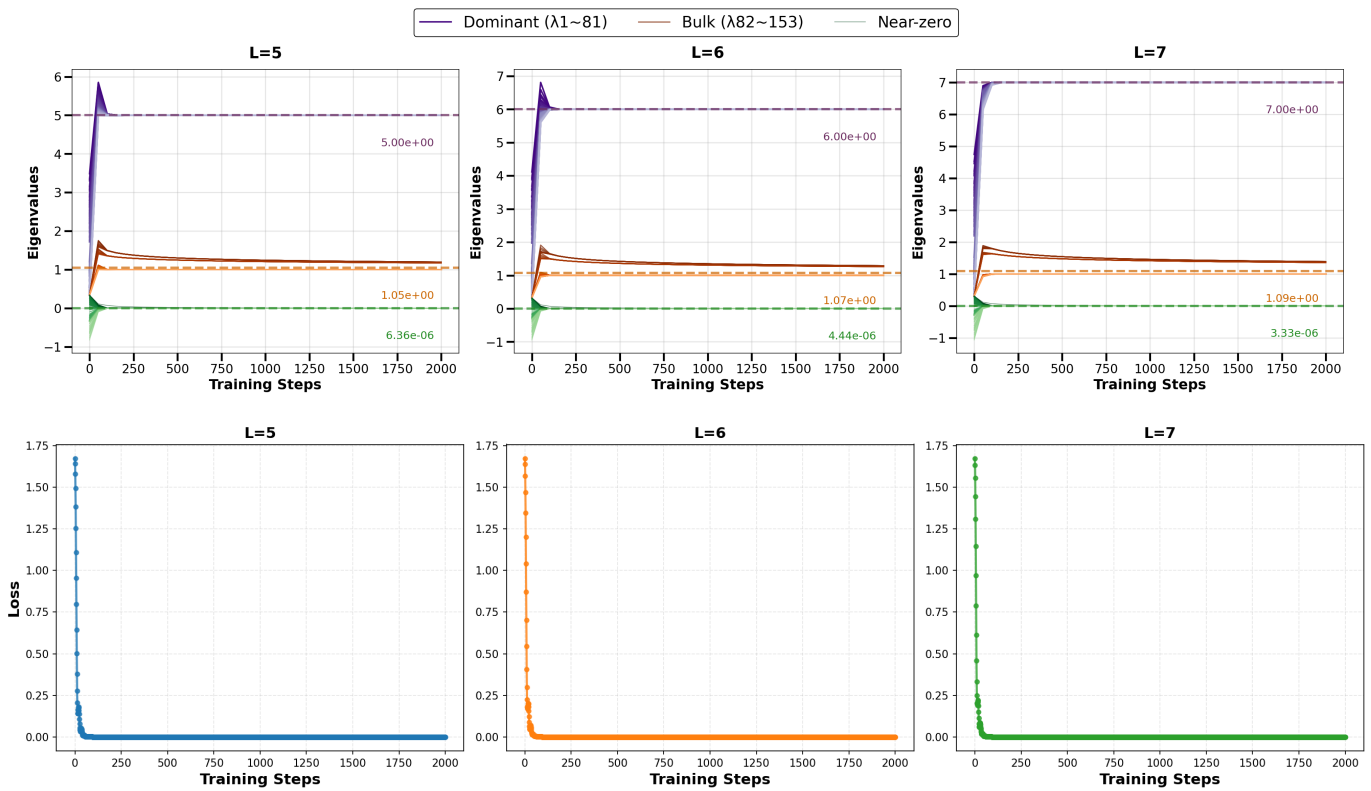


Fig. 15: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

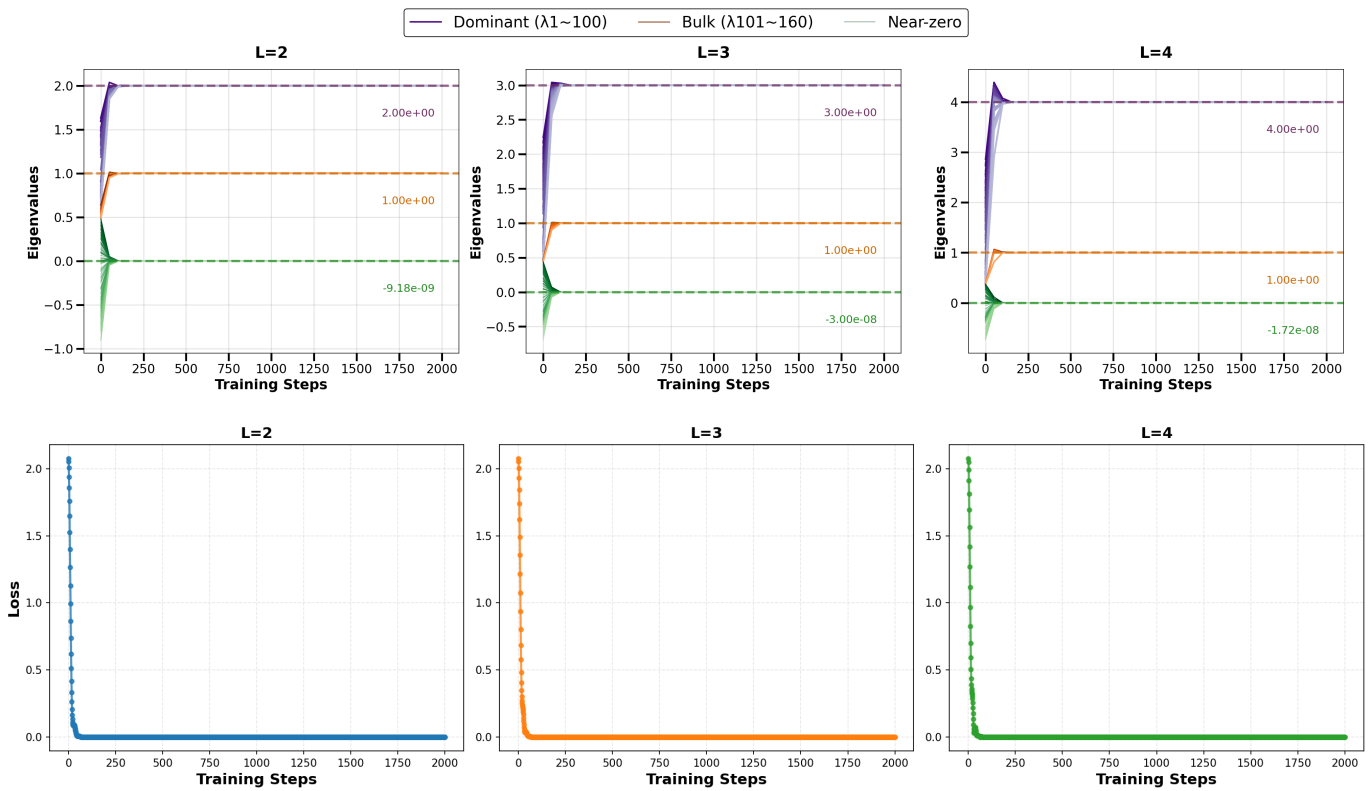


Fig. 16: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).

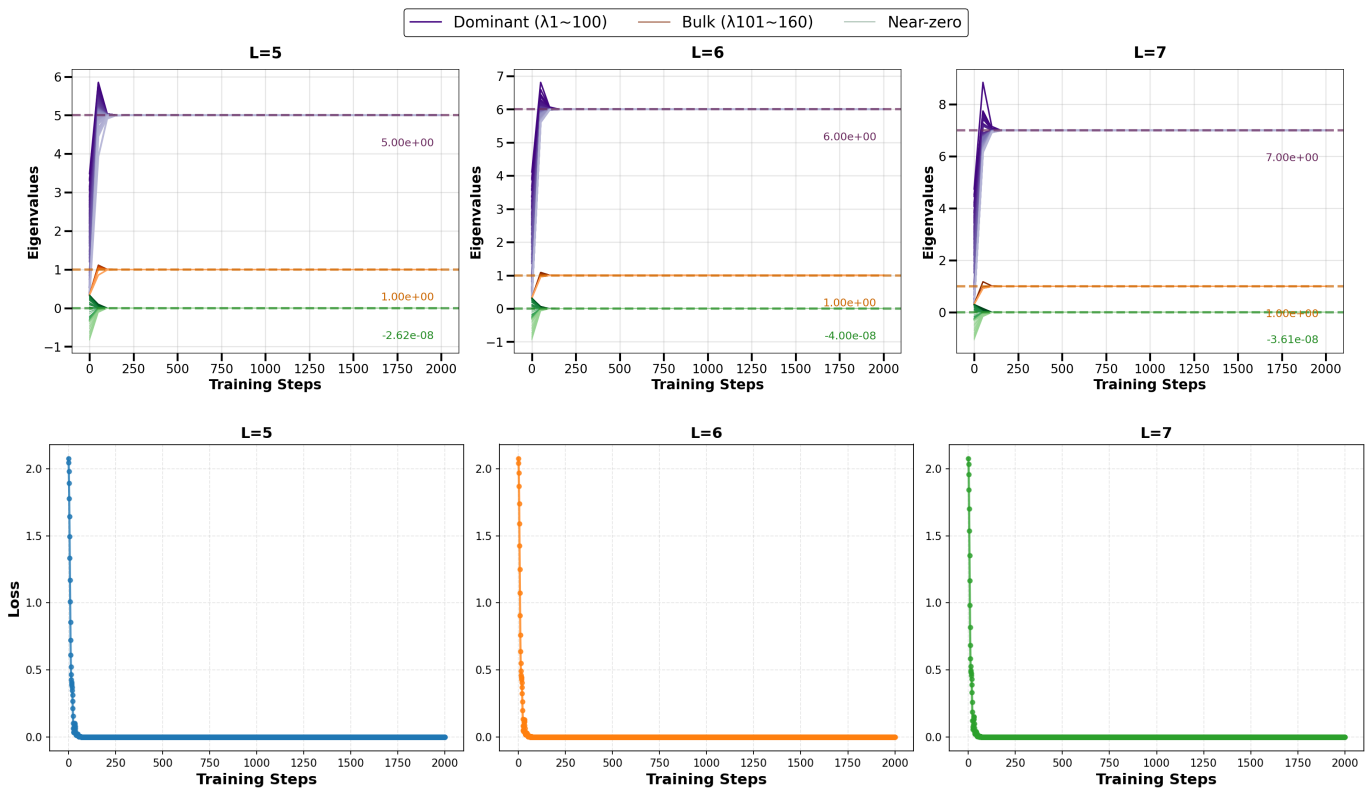


Fig. 17: **Eigenvalue evolution.** The curves are color-coded by subspace: purple for the dominant space, orange for the bulk space, and green for the near-zero space. The dominant space has a dimension of rank^2 , while the combined dimension of the dominant and bulk spaces equals the product of the input and output dimensions. The final eigenvalues of the dominant space converge to L times those of the bulk space. The panels correspond to $L = 3$ (left), $L = 4$ (middle), and $L = 5$ (right).