POINTWISE GENERALIZATION IN DEEP NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

We address the long-standing question of why deep neural networks generalize by establishing a complete pointwise generalization theory for fully connected networks. For each trained model, we equip the hypothesis with a pointwise Riemannian dimension through the effective ranks of the *learned* feature matrices across layers, and derive hypothesis- and data-dependent generalization bounds. These spectrum-aware bounds break long-standing barriers and are orders of magnitude tighter in theory and experiment, rigorously surpassing bounds based on model size, products of norms, and infinite-width linearizations. Analytically, we identify structural properties and mathematical principles that explain the tractability of deep nets. Empirically, the pointwise Riemannian dimension exhibits substantial dimensionality reduction, decreases with increased over-parameterization, and captures feature learning and the implicit bias of optimizers across standard datasets and modern architectures. Taken together, these results show that deep networks are mathematically tractable in the practical regime and that their generalization is sharply explained by pointwise, spectrum-aware complexity.

1 Introduction

Deep learning has ushered in a new era of AI, delivering striking generalization across a wide range of scientific tasks. Yet these successes are predominantly empirical; theory has not kept pace. Paradoxically, despite massive overparameterization, especially for large language models, classical theory predicts severe overfitting, whereas practice shows strong generalization. The resulting gap has fueled a prevailing pessimism that neural networks are opaque "black boxes" resistant to principled explanation. This paper closes a key part of that gap by resolving the generalization problem for the canonical neural network—fully connected deep neural network (DNN). Under minimal, verifiable spectral conditions, we prove that fully connected deep networks fall into the tractable family—on a rigorous footing comparable to sparse linear models and low-rank matrix factorization—rather than the unconstrained "general" overparameterized class. To our knowledge, this is the first rigorous account in decades that establishes generalization in fully connected deep networks as tractable by the accepted standards of learning theory.

We study standard fully connected (feed-forward) networks on a dataset $X = [x_1, \ldots, x_n] \in \mathbb{R}^{d_0 \times n}$, where each column is one input example. The network has widths d_1, \ldots, d_L , and weight matrices $W_l \in \mathbb{R}^{d_l \times d_{l-1}}$ for $l = 1, \ldots, L$. We define the *feature matrix* at layer l by the recursion

$$F_l(W,X) := \sigma_l(W_l F_{l-1}(W,X)) \in \mathbb{R}^{d_l \times n}, \qquad l = 1,\dots, L, \tag{1}$$

where $F_0 := X$ and the nonlinear activation σ_l acts columnwise. Each *column* of F_l is the feature vector of one data point at layer l; each *row* of F_l is the activation of one neuron across the dataset.

Our focus is the *generalization gap*—the difference between test and training loss at the learned weights W. Informally (up to universal constants and mild logarithms), we prove that this gap is controlled by the *effective dimension* of the learned features: uniformly over every $W \in \mathbb{R}^{\sum_l d_l \cdot d_{l-1}}$,

$$\mathcal{L}_{\text{test}}(W) - \mathcal{L}_{\text{train}}(W) \lesssim \sqrt{\frac{1}{n} \sum_{l=1}^{L} (d_l + d_{l-1}) d_{\text{eff}} \Big(F_{l-1}(W, X) F_{l-1}(W, X)^{\top} \Big)}.$$
 (2)

Here $d_{\text{eff}}(\cdot)$ denotes the (layerwise) *effective dimension*—a smoothed notion of rank—of the feature Gram matrix $F_{l-1}(W,X)F_{l-1}(W,X)^{\top}$, i.e., the number of meaningful directions the feature data actually occupies at that layer. Intuitively, each layer contributes a term proportional to its size (d_l+d_{l-1}) multiplied by how many directions its features $F_{l-1}(W,X)$ truly use, d_{eff} . When features are correlated, low rank, or exhibit a rapidly decaying spectrum (a few large eigenvalues dominating many small ones), d_{eff} is small, so the bound remains tight even for very wide/deep networks and improves with sample size as $1/\sqrt{n}$. Such "feature compression" phenomena is widely observed in modern deep learning (Huh et al., 2021; Wang et al., 2025; Parker et al., 2023). Strikingly, in our experiments, increasing overparameterization often induces pronounced *feature-rank compression*: the bound (2) decreases as model size grows (Section A); for example, in ResNet trained on CIFAR-10, a majority of layers compress to (near-)zero effective rank.

Inequality (2) yields a strong *uniform, hypothesis-dependent* guarantee, which we term *pointwise* generalization. It tracks how features evolve across layers of the trained model and explains overparameterization in practice. Moreover, the right-hand side of (2) can be used directly as a regularizer, leading to algorithms that adapt to the effective ranks around a benchmark W^* (Section 4.2). Under minimal spectral conditions, our theory places fully connected networks in the same complexity class as sparse linear models and low-rank matrix factorization: generalization is governed by a effective dimension rather than parameter count. The spectrum-aware effective-dimension notion we adopt is standard and minimax-sharp in linear and kernel settings (Even & Massoulié, 2021). In contrast, existing bounds either rely on infinite-width linearizations (e.g., NTK), are exponential in products of spectral norms (from (Bartlett & Mendelson, 2002) to (Bartlett et al., 2017)), or scale with model size (e.g., VC dimension (Bartlett et al., 2019)). Our bounds avoid these pathologies and provide a pointwise, spectral account that matches known upper and lower rates, thereby establishing fully connected deep networks as tractable by the accepted standards of learning theory.

1.1 Contributions

The paper is organized into three parts—(i) a pointwise generalization framework (Section 2, (ii) structural principles of deep networks (Section 3-4), and (iii) empirical validation (Section A). We elaborate noveltes in each.

Pointwise Generalization and Finite-Scale Geometry. Classical tools (e.g., Rademacher complexity, uniform covers, products of norms) measure global complexity and are often too coarse for modern deep nets: they miss how a specific trained model uses its learned features across layers. We propose a pointwise framework that targets the model actually trained and yields bounds with matching upper and lower rates via a finite-scale notion of pointwise dimension—achieving the precision of generic chaining—while assigning each hypothesis a pointwise dimension that governs its error. This yields a geometric view of generalization: a *finite-scale*, spectrum-aware geometry driven by dimension reduction (as opposed to infinitesimal limits), which clarifies the nature of generalization and the sources of its difficulty. Our bounds can also be read as optimally analyzed PAC–Bayes: the analysis admits data-dependent priors and deterministic predictors while retaining posterior adaptivity, thereby overcoming core limitations of standard PAC–Bayes approaches.

Structual Principles and Tight Bounds for Neural Networks. We develop a non-perturbative approach that uses exact telescoping decompositions (rather than Taylor linearizations) to preserve the finite-scale geometry of deep networks. This yields our first structural principle: cross-layer correlations factor through the feature matrices and approximately preserve a pointwise linear structure. We then show that bounding the pointwise dimension reduces to the gold standard of effective dimension on local charts, and we extend this to a global statement by constructing an ellipsoidal covering over the set of subspaces (Grassmannian). This extension—novel beyond the classical differential—geometric/Lie—algebraic treatments—establishes our second structural principle: the complexity of the global atlas (covering reference top eigenspaces) does not exceed that of the local charts. Building on these principles, we introduce a pointwise Riemannian dimension—a spectrum-aware, layerwise effective complexity—that governs generalization at the trained model and yields tight, analyzable bounds. We review each step and argue that the resulting bounds are tight in a qualified sense; moreover, they unconditionally reduce to spectral—norm bounds (see Appendix E.3).

Empirical Findings and Evidences. The experiments are designed to systematically examine three central questions in modern deep learning: (i) why does overparameterization often improve gener-

alization? (ii) how does feature learning evolve during training? and (iii) what implicit regularization is encoded by the baseline optimizer? Across the experimental results, we observe that (i) the overparameterization impressively leads to decreasing Riemannian Dimension; (ii) feature learning compresses the effective ranks of learned features during the training; and (iii) SGD with momentum implicitly regularizes the Riemannian Dimension.

2 THE NATURE OF GENERALIZATION

In this section, we develop our pointwise framework for generalization analysis, which introduces a tight tool-pointwise dimension-to characterize generalization. We illustrate its advancement to existing methodologies and bring some new understandings on the nature of generalization.

2.1 Pointwise Generalization as BEST PAC-Bayes Optimization

Let \mathcal{F} be a hypothesis class, z be random data drawn from an unknown distribution \mathbb{P} (e.g., inputlabel pair z=(x,y)), and $\ell(f;z)$ be real-valued loss function. Denote by \mathbb{P}_n the empirical distribution supported on an i.i.d. sample $\{z_i\}_{i=1}^n \sim \mathbb{P}^n$. Our goal is to control the *generalization gap* $(\mathbb{P}-\mathbb{P}_n)\ell(f;z)$ in the following manner: for $\delta \in (0,1)$, with probability at least $1-\delta$, uniformly over every $f \in \mathcal{F}$,

$$(\mathbb{P} - \mathbb{P}_n)\ell(f; z) := \mathbb{E}_{z \sim \mathbb{P}} \left[\ell(f; z) \right] - \frac{1}{n} \sum_{i=1}^n \ell(f; z_i) \le C \sqrt{\frac{d(f) + \log \frac{1}{\delta}}{n}}, \tag{3}$$

where d(f) is a hypothesis-dependent complexity measure that aims to characterize the intrinsic complexity of every *trained* hypothesis f, different from canonical uniform convergence analysis.

In the spirit of (3) we introduce the core concept in this section—pointwise dimension, a concept strengthen several established generalization methodologies such as PAC-Bayesian analysis, Kolmogorov complexity, generic chaining, and we then illustrate its tightness in characterizing the generalization by two theorems. Throughout the paper, we use the term "metric" ϱ in the pseudometric sense: it satisfies all the usual metric axioms except that $\varrho(f_1, f_2) = 0$ need not imply $f_1 = f_2$.

Definition 1 (Pointwise Dimension) Given a function class \mathcal{F} , a metric ρ on \mathcal{F} , and a data-dependent prior π over \mathcal{F} , the local dimension at f with scale ε is defined as the log inverse density of the ε - ball $B_{\varrho}(f,\varepsilon) = \{f' \in \mathcal{F} : \varrho(f,f') \leq \varepsilon\}$ centered at f:

$$\log \frac{1}{\pi(B_{\varrho}(f,\varepsilon))}. (4)$$

To prepare for our first theorem, we define the loss-induced empirical $L_2(\mathbb{P}_n)$ metric $\varrho_{n,\ell}$ as

$$\varrho_{n,\ell}(f_1, f_2) = \sqrt{\mathbb{E}_{z \sim \mathbb{P}_n} \left[(\ell(f_1; z) - \ell(f_2; z))^2 \right]}.$$

Theorem 1 (One-Shot Bound) *Let* π *be any data-dependent prior on a function class* \mathcal{F} , and loss $\ell(f;z)$ bounded by [0,1]. Then for every $\delta \in (0,1)$, with probability at least $1-\delta$ over n i.i.d. draws $z_1, \ldots, z_n \sim \mathbb{P}$, uniformly over all $f \in \mathcal{F}$,

$$(\mathbb{P} - \mathbb{P}_n)\ell(f;z) \leq \inf_{\varepsilon > \sqrt{1/n}} \left\{ 2\varepsilon + \sqrt{\frac{2\log\frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\varepsilon))}}{n}} \right\} + \sqrt{\frac{4}{n}} + 3\sqrt{\frac{\log\frac{1}{\delta}}{2n}}.$$

Intuitively, pointwise dimension not only concerns the prior mass on a single hypothesis but works for general, incountable \mathcal{F} by taking the piror mass over a ball centered at f. This overcomes key limitations of previous hypothesis-dependent bounds such as Occam's razor bound and Kolmogorov complexity (Lotfi et al., 2022). Additionally, our perspective brings the best possible PAC-Bayesian mechanism, where the PAC-Bayes is seen as a promising tool for sharpening generalization. The proof develops the generalization error as an bias-variance optimization problem on the user-chosen posterior. Through very novel techniques such as symmetrization and conditioning, we permits data-dependent priors (not allowed in classical PAC-Bayes) and makes the framework applicable to every

single hypothesis beyond artificial randomization. This theorem along, with proved optimality of chosen posterior and closed-form final expression, greatly improves relevant works in the area such as (Hinton & Van Camp, 1993) and (Dziugaite & Roy, 2017).

We present a second bound that strengthens the above one-shot bound via a multi-scale integral (known as generic chaining Talagrand (2005)). It covers rich classes whose pointwise dimension can grow as $O(d(f) \varepsilon^{-2})$ while still achieving a generalization rate of $\sqrt{d(f)/n}$.

Theorem 2 (Generic Chaining Bound and Global Lower Bound) For loss $\ell(f;z)$ bounded in [0,1], (i) there exists an absolute constant C>0 such that for any data-dependent prior π on \mathcal{H} and any $\delta\in(0,1)$, with probability at least $1-\delta$, uniformly over every $f\in\mathcal{F}$

$$(\mathbb{P} - \mathbb{P}_n)\ell(f; z) \le C \left(\frac{1}{\sqrt{n}} \int_0^\infty \sqrt{\log\left(\frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\varepsilon))}\right)} d\varepsilon + \sqrt{\frac{\log\frac{\log(2n)}{\delta}}{n}} \right);$$

(ii) under mild measurability conditions there are absolute constants c, c' > 0 so that

$$\mathbb{E}\left[\inf_{\pi}\sup_{f\in\mathcal{F}}\left((\mathbb{P}-\mathbb{P}_n)\ell(f;z)-\frac{c}{\sqrt{n\log n}}\int_0^{\infty}\sqrt{\log\frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\varepsilon))}}d\varepsilon\right)+\frac{c'\sup_{\mathcal{F}}\mathbb{E}[\ell(f;z)]}{\sqrt{n\log n}}\right]\geq 0,$$

where notation \mathbb{E} means taking expectation over samples.

The integral upper bound in Theorem 2 is *tight* in the following sense: no uniform improvement (valid simultaneously for all hypotheses and priors) is possible. This is witnessed by a matching lower bound. A strictly *pointwise* lower bound (depending on the realized hypothesis) is generally unattainable, because the prior π must be hypothesis—blind (a "no free lunch" constraint). Theorem 2 extends Talagrand's celebrated generic chaining to *pointwise* generalization bounds. Consequently, it is fundamentally stronger than classical entropy—integral bounds based on global covering numbers, e.g., Dudley's integral (Vershynin, 2009), whose integrand takes a supremum over the entire class \mathcal{F} rather than localizing at the realized hypothesis.

We defer technical innovations and connections to existing methodologies to the appendix. The main takeaway is that the proposed *pointwise dimension* is a powerful and precise descriptor: it tightly characterizes pointwise generalization.

2.2 GENERALIZATION IS FINITE-SCALE DIMENSION

We advocate the viewpoint that the nature of generalization's is a *finite-scale* notion of dimension. Concretely, our proposed pointwise dimension 4 is evaluated at a finite resolution—capturing the model's intrinsic, spectrum-aware complexity and drives dimension reduction at this scale. This stands in sharp contrast to infinitesimal-scale geometric notions (which reduce to model size), and therefore miss the structure that governs predictive performance.

Asymptotic vs. finite-scale dimension. A central notion to geometry is asymptotic pointwise dimension, and a classical definition is

$$\lim_{\varepsilon \to 0} \frac{\log \pi(B(f,\varepsilon))}{\log \varepsilon},$$

which is essential to fractal geometry (e.g., Chapter 10.1 in (Falconer, 1997)) and Riemannian geometry (Lutz, 2016), e.g., in the classical characterizations of Hausdorff and packing dimensions. According to this definition, geometric dimension is inherently infinitesimal: it studies limit behavior $\varepsilon \to 0$ at the point f. A key point that distinguishes generalization to geometry is that generalization studies the finite-scale dimension; and our pointwise dimension $\log \frac{1}{\pi(B_\varrho(f,\varepsilon))}$ clearly reduces as ε increases; thus finite-scale study of geometry leads to significant dimension reduction. In Theorem 1, the goal of generalization is to identify the best finite scale ($\varepsilon^* \approx$ resulted bound), and at this scale our pointwise dimension (4) could be much smaller than the asymptotic dimensions, which allows tractable generalization in overparameterized models.

3 DEEP NEURAL NETWORKS AND RIEMANNIAN DIMENSION

We consider fully connected (feed-forward) networks that map an input $x \in \mathbb{R}^{d_0}$ to an output $f_L(W,x) \in \mathbb{R}^{d_L}$. The architecture is specified by widths d_0,\ldots,d_L and weight matrices $W=\{W_1,\ldots,W_L\}$ with $W_l \in \mathbb{R}^{d_l \times d_{l-1}}$ for $l=1,\ldots,L$. Let σ_1,\ldots,σ_L be nonlinear activations (e.g., ReLU), acting componentwise on column vectors, and each $\sigma_l:\mathbb{R}^{d_l}\to\mathbb{R}^{d_l}$ is assumed 1-Lipschitz. The network's forward map is the composition

$$f_L(W,x) := \sigma_L \Big(W_L \, \sigma_{L-1} \big(W_{L-1} \, \cdots \, \sigma_1(W_1 x) \big) \Big).$$

Let $X = [x_1, ..., x_n] \in \mathbb{R}^{d_0 \times n}$ collect the n training inputs as columns. For each layer $l \in \{1, ..., L\}$, define the depth-l map and the corresponding feature matrix

$$f_l(W,x) := \sigma_l(W_l \sigma_{l-1}(W_{l-1} \cdots \sigma_1(W_1x))), F_l(W,X) := [f_l(W,x_1) \cdots f_l(W,x_n)] \in \mathbb{R}^{d_l \times n}.$$

Equivalently (full, non-recursive form consist with (1)),

$$F_l(W,X) = \sigma_l(W_l \sigma_{l-1}(W_{l-1} \cdots \sigma_1(W_1X))),$$

where for a matrix $A = [a_1, \ldots, a_n]$ we write $\sigma_l(A) := [\sigma_l(a_1), \ldots, \sigma_l(a_n)]$. Thus $F_L(W, X)$ collects the network outputs on the dataset X.

We use $\|\cdot\|_F$ for the Frobenius norm and $\|\cdot\|_{op}$ for the spectral norm. The empirical $L_2(\mathbb{P}_n)$ distance between two hypotheses W, W' is (a $1/\sqrt{n}$ scaling is used to keep consistency with Section 2)

$$\varrho_n(W, W') := \sqrt{\|F_L(W, X) - F_L(W', X)\|_F^2/n}.$$

The function-level empirical metric and generalization statements in Section 2 for the loss $x \mapsto \ell(f_L(W,x),y)$ at data-label pairs z=(x,y) specialize, on the dataset X, to the metric ϱ_n defined above. We assume the loss $\ell(\cdot,y)$ is β -Lipschitz in its first argument with respect to $f_L(W,x)$, which bridges the metric $\varrho_{n,\ell}$ studied in Section 2 to ϱ_n defined on the weight space.

3.1 Non-Perturbative Expansion and Layer-wise Correlation

Throughout, our finite-scale analysis relies on *non-perturbative* expansions. Borrowing terminology from theoretical physics, "non-perturbative" here means we avoid Taylor/derivative expansions and instead use exact, telescoping algebraic identities that hold at finite scale. For example,

$$W_2'W_1' - W_2W_1 = W_2'(W_1' - W_1) + (W_2' - W_2)W_1, \qquad {\Sigma'}^{-1} - {\Sigma'}^{-1} = {\Sigma'}^{-1}(\Sigma - \Sigma')\Sigma^{-1},$$

with analogous decompositions used throughout. This viewpoint preserves the full finite-scale geometry of deep networks, rather than linearizing around an infinitesimal neighborhood.

To present our non-perturbative expansion for DNNs, we define *local Lipchitz constant* $M_{l\to L}(W,\varepsilon)$, which characterizes the sensitivity of the layer L output, F_L , to variations in layer l's output, within a neighborhood around F_l . Formally, we assume that for every $W' \in B_{\rho_n}(W,\varepsilon)$

$$||F_L(F_l(W',X),\{W_i'\}_{i=l+1}^L) - F_L(F_l(W,X),\{W_i'\}_{i=l+1}^L)||_{\mathbf{F}} \le M_{l\to L}(W,\varepsilon)||F_l(W',X) - F_l(W,X)||_{\mathbf{F}}.$$

Local Lipschitz constants are typically much smaller than products of spectral norms and can be computed by formal-verification toolchains (Shi et al., 2022). In our bounds these constants appear only inside *logarithmic factors*, so they do not affect the leading rates. For completeness, we discuss them carefully in the appendix. We propose a telescoping decomposition to replace conventional Taylor expansion, where in each summand the only difference lie in W_l and W_l .

$$\begin{split} &F_L(W',X) - F_L(W,X) \\ &= \sum_{l=1}^L [\underbrace{\sigma_L(W'_L \cdots W'_{l+1})}_{\text{controled by} M_{l \to L}} \underbrace{\sigma_l}_{\text{by1}} (W'_l \underbrace{F_{l-1}(W,X)}_{\text{learned feature}})) - \sigma_L(W'_L \cdots W'_{l+1} \sigma_l(W_l \underbrace{F_{l-1}(W,X)}_{\text{learned feature}}))], \end{split}$$

Note that this is a *non-perturbative* expansion that holds unconditionally and does not rely on infinitesimal approximation, and crucially keeps the *learned* feature $F_{l-1}(W, X)$ at the *trained* weight W. From this decomposition and applying basic inequalities, we have the following key lemma.

Lemma 1 (Non-Perturbative Feature Expansion) For all $W' \in B_{\rho_n}(W, \varepsilon)$,

$$||F(W',X) - F(W,X)||_{\mathbf{F}}^2 \le \sum_{l=1}^{L} L \cdot M_{l \to L}[W,\varepsilon]^2 \cdot ||(W'_l - W_l)F_{l-1}(W,X)||_{\mathbf{F}}^2.$$
 (5)

The lemma captures the first structural principle of fully connected DNNs: cross-layer correlations mostly pass through the feature matrices, preserving an approximate pointwise linear structure.

Since enlarging the metric only shrinks metric balls and hence *increases* the pointwise dimension (4) we analyze in Section 2, it suffices to analyze pointwise dimension under the *pointwise ellipsoidal metric* that appears on the right-hand side of Lemma 1. Concretely, $G_l(W) := F_{l-1}(W,X) F_{l-1}(W,X)^{\top}$, the feature Gram matrix from layer l-1, faithfully encodes the spectral information induced by the network–data geometry. Working with the corresponding pointwise ellipsoidal metric yields sharp, *pointwise, spectrum-aware* bounds with the desired properties for deep networks, and underpins our tractability results (with the structural principles and technical innovations to developed in the next subsection).

3.2 HIERARCHICAL COVERING FROM LOCAL CHART TO GLOBAL ATLAS

Lemma 1 suggests that the following *pointwise ellipsoidal metric* dominates $n \cdot \varrho_n$ at every W (here, NP stands for "non-perturbative"):

$$G_{\text{NP}}(W) = \text{blockdiag}\left(\cdots, LM_{l \to L}^{2}(W, \epsilon) \cdot F_{l-1}(W, X) F_{l-1}^{\top}(W, X) \otimes I_{d_{l}}, \cdots\right)$$

$$\varrho_{G_{\text{NP}}(W)}(W, W') = \text{vec}(W' - W)^{\top} G_{\text{NP}}(W) \text{vec}(W' - W). \tag{6}$$

We are therefore interested in bounding the enlarged pointwise dimension under the pointwise ellipsoidal metric $\varrho_{G_{\mathrm{NP}}(W)}$: $\log \frac{1}{\pi(B_{\varrho_{n}}(f(W,\cdot),\epsilon))} \leq \log \frac{1}{\pi(B_{\varrho_{G_{\mathrm{NP}}}(W)}(W,\sqrt{n}\epsilon))}$.

Golden standard: effective dimension. Classical studies of static ellipsoidal metrics suggest that if π is chosen to be uniform over the top-r eigenspace of $G_{\rm NP}(W)$, then one can achieve a tight effective dimension as follows: define

$$r_{\text{eff}}(G(W), R, \varepsilon) := \max\{k : \lambda_k(G(W), R^2 \ge n\varepsilon^2/2\},$$
 (7)

and spectrum-aware effective dimension

$$d_{\text{eff}}(G(W), R, \varepsilon) := \frac{1}{2} \sum_{k=1}^{r_{\text{eff}}(G(W), R, \varepsilon)} \log \left(\frac{8R^2 \lambda_k}{n\varepsilon^2} \right)$$
 (8)

This definition serves as a gold standard for static ellipsoidal metrics and is asymptotically tight, as established by the covering number of the unit ball with ellipsoids in Dumer et al. (2004). For brevity, we write r for $r_{\rm eff}(G(W),R,\varepsilon)$, and denote by $\mathcal{V}\subset\mathbb{R}^p$ the r-dimensional subspace corresponding to the top- $r_{\rm eff}$ eigenspace of G(W).

Key challenge: prior independence from W. However, the main challenge is that the construction of π cannot rely on knowledge of W, including its top- $r_{\rm eff}$ eigenspace, yet still capture the underlying geometric structure. The next lemma extends classical results on static ellipsoidal metrics by showing that a uniform prior over a reference subspace $\bar{\mathcal{V}}$ suffices to bound the pointwise dimension for all W whose top-r eigenspace of G(W) can be approximated by $\bar{\mathcal{V}}$.

Lemma 2 (Pointwise Dimension via Reference Subspace) Consider the weight space $B_2(R) \subset \mathbb{R}^p$ for vectorized weights, and a pointwise ellipsoidal metric defined via PSD G(W). Let $\overline{\mathcal{V}} \subseteq \mathbb{R}^p$ be an r-dimensional subspace. Define the prior $\pi_{\overline{\mathcal{V}}} = \text{Unif}(B_2(1.58R) \cap \overline{\mathcal{V}})$. Then, uniformly over all (W, ε) such that the top-r eigenspace \mathcal{V} of G(W) can be approximated by $\overline{\mathcal{V}}$ to precision

$$\varrho_{\operatorname{proj},G(W)}(\mathcal{V},\bar{\mathcal{V}}) := \|G(W)^{1/2} (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\| \le \frac{\sqrt{n\varepsilon}}{4R},\tag{9}$$

we have

$$\log \frac{1}{\pi_{\bar{\mathcal{V}}}(B_{\varrho_{G(W)}}(W,\sqrt{n}\varepsilon))} \, \leq \, \frac{1}{2} \sum_{k=1}^{r_{\mathrm{eff}}(G(W),R,\varepsilon)} \log \left(\frac{40R^2\lambda_k}{n(\varepsilon^2-\varepsilon_1^2)} \right) \, = \, d_{\mathrm{eff}}(\Sigma,\sqrt{5}R,\varepsilon).$$

In (9), $\mathcal{P}_{\mathcal{V}}$ denotes the orthogonal projector onto the subspace \mathcal{V} , and $\varrho_{\text{proj},G(W)}$ thus defines an ellipsoidal projection metric between subspaces. Further details are provided in the appendix.

Hierarchical covering (mixture prior over subspaces). We employ a hierarchical covering argument. For each reference subspace $\bar{\mathcal{V}}$, the bottom-level prior (uniform on $\bar{\mathcal{V}}$) can achieve a tight pointwise dimension bound for all "local" weights W whose top-r eigenspace of G(W) can be well-approximated by $\bar{\mathcal{V}}$. At the top level, we then construct a prior over $\bar{\mathcal{V}}$. By combining these two levels of priors, we obtain a pointwise dimension bound using a prior π that is completely blind to the choice of W. To formalize this, we introduce a top-level distribution μ over the Grassmannian

$$Gr(p,r) := \{r - \text{dimensional linear subspaces of } \mathbb{R}^p \}$$

the collection of all r-dimensional subspaces, and define

$$\pi(W) = \sum_{\mathcal{V}} \mu(\mathcal{V}) \, \pi_{\mathcal{V}}(W).$$

We refer to this two-stage construction as the hierarchical covering argument. Under the resulting prior π , the following bound holds uniformly over all $W \in \mathcal{W}$, the pointwise dimension $\log \frac{1}{\pi(B_{\varrho_G(W)}(W,\sqrt{n}\epsilon))}$ is bounded by two parts:

$$\underbrace{\log \frac{1}{\mu(B_{\varrho_{\text{proj}},G(W)}(\mathcal{V},\sqrt{n\epsilon/4R}))}}_{\text{covering Grassmannian (global atlas)}} + \underbrace{\sup_{\bar{\mathcal{V}}\in B_{\varrho_{\text{proj}},G(W)}(\mathcal{V},\sqrt{n\epsilon/4R})} \log \frac{1}{\pi_{\bar{\mathcal{V}}}(B_{\varrho_{G(W)}}(W,\sqrt{n\epsilon}))}}_{\text{covering local charts}}, \quad (10)$$

In differential–geometric terms, our argument has two components. Local (chart) analysis: fixing a reference subspace $\bar{\mathcal{V}}$, we use effective dimension as the gold standard to determine the metric entropy of the corresponding local chart. Global (atlas) covering: we cover the Grassmannian by such reference subspaces, i.e., we bound the metric entropy of the global atlas and account for the cost of transitioning across charts. Lemma 2 controls the local part, while the following new result on the ellipsoidal Grassmannian controls the global part:

Lemma 3 (Ellipsoidal Covering of the Grassmannian manifold) Consider the Grassmannian $\operatorname{Gr}(d_l, r_l)$. For uniform prior $\mu = \operatorname{Unif}(\operatorname{Gr}(d_l, r_l))$, we have that for every $\mathcal{V} \in \operatorname{Gr}(d_l, r_l)$, every $\epsilon > 0$ and every PSD matrix Σ with eigenvalues $\lambda_1 \geq \cdots \lambda_{d_l}$, we have the pointwise dimension bound

$$\log \frac{1}{\mu(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V},\epsilon))} \leq \frac{1}{2}(d_l - r_l) \sum_{k=1}^{r_l} \log \frac{C \max\{\lambda_k,\epsilon^2/42\}}{\epsilon^2} + r_l \sum_{k=1}^{d_l - r_l} \log \frac{C \max\{\lambda_k,\epsilon^2/42\}}{\epsilon^2},$$

where C > 0 is an absolute constant.

The result above is mathematically significant in its own right. It extends the classical metric-entropy (covering number) theory for the Grassmannian—where log covering number $\asymp r(d-r)\log(C/\epsilon)$ under the *isotropic* projection metric— to an *ellipsoidal* (anisotropic) metric that captures feature—and model—induced geometry. This generalization translates the traditional differential-geometric and Lie-algebraic treatments (see Appendix D) and, we believe, illustrates a two—way exchange: deep mathematical structure is essential to understanding generalization in modern neural networks, and, conversely, generalization theory can motivate new questions and results in pure mathematics.

Leveraging the block–decomposable structure in (6), we obtain:

Theorem 3 (Riemannian Dimension for DNN) Consider the weight space $B_{\mathbf{F}}(R)$, and a pointwise ellipsoidal metric defined via the ellipsoidal metric $G_{NP}(W)$ defined in (6). Define the pointwise Riemannian dimension

$$d_R(W, \epsilon) = \sum_{l=1}^{L} \left(\underbrace{d_l \cdot d_{\text{eff}}(A_l(W))}_{\text{covering local charts}} + \underbrace{d_{l-1} \cdot d_{\text{eff}}(A_l(W))}_{\text{covering global atlas}} + \underbrace{\log(d_{l-1}n)}_{\text{covering discrete}} \right),$$

where $d_{\text{eff}}(A_l(W))$ is abbreviation of $d_{\text{eff}}(A_l(W), C \max\{\|W\|_{\mathbf{F}}, R/2^n\}, \epsilon)$ and C is a positive absolute constant. Then we have the pointwise dimension bound: uniformly over all $W \in B_{\mathbf{F}}(R)$,

$$\log \frac{1}{\pi(B_{\varrho_n}(f(W,\cdot),\epsilon))} \le d_R(W,\epsilon).$$

This concludes our program for fully connected networks: we establish *Riemannian dimension* as a principled complexity measure that explains—and sharply bounds—generalization. We summarize the *second structural principle of fully connected DNNs*: The complexity of the *global atlas* (covering the space of reference top eigenspaces) remains commensurate with the layerwise, spectrum–aware complexity of covering the *local charts*.

4 GENERALIZATION BOUNDS AND COMPARISON

4.1 GENERALIZATION BOUND FOR DNNs

We are now ready to state our generalization bound for fully connected DNN here. Combining Theorem 3 and Theorem 2, we have

Theorem 4 (Generalization Bounds for DNNs) Let the loss $\ell(f(W,x),y)$ be bounded in [0,1] and β -Lipchitz with respect to f(W,x), for every $\delta \in (0,1)$, with probability at least $1-\delta$, uniformly over all $W \in \mathcal{W}$,

$$(\mathbb{P} - \mathbb{P}_n)\ell(f(W, x), y) \le C_1 \left(\frac{\beta}{\sqrt{n}} \int_0^\infty \sqrt{d_R(W, \epsilon)} d\epsilon + \sqrt{\frac{\log \frac{\log(2n)}{\delta}}{n}}\right),$$

where the admitted Riemannian Dimension is

$$d_R(W,\epsilon)$$

$$= \sum_{l=1}^{L} \left((d_l + d_{l-1}) \sum_{k=1}^{r_{\text{eff}}[W,l]} \log \frac{8C_2^2 \lambda_k(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_F^2, R^2/4^n\} L M_{l \to L}^2(W, \epsilon)}{n\epsilon^2} + \log(d_{l-1}n) \right),$$

where F_{l-1} is learned feature $F_{l-1}(W,X)$; $r_{\text{eff}}[W,l] := r_{\text{eff}}(LM_{l \to L}^2(W,\epsilon)F_{l-1}F_{l-1}^\top,C_2\max\{||W||_F,R/2^n\},\epsilon)$ is abbreviation of effective rank; and C_1,C_2 are positive absolute constant.

We conclude by reviewing our comprehensive theory for generalization in fully connected networks and justifying the tightness of the resulting bounds. **First**, in Section 2 we develop a framework based on *pointwise dimension*. The upper and lower bounds match in a qualified (non-uniform) sense (see remarks after Theorem 2), and the framework has a profound connection to finite-scale geometry—evidence that this is the right organizing principle. **Second**, Section 3 introduces a *non-perturbative* expansion. Lemma 1 applies Cauchy—Schwarz layerwise (treating each layer as a block). While there may be room to improve depth dependence, the telescoping decomposition (5) is an *equality*, so the expansion is generally sharp (and fully avoid linearization). **Third**, the hierarchical covering argument shows that the resulting *Riemannian dimension* bound matches the gold standard of *effective dimension*. Thus our pointwise, spectrum-aware bounds achieve the optimal form dictated by static ellipsoid theory, now in strongly correlated deep networks.

4.2 IMPLICIT BIAS AND ALGORITHMIC IMPLICATION

Pointwise Dimension is Implicit Bias. A central tenet in deep learning generalization is implicit bias (regularization favored by the algorithms) (Vardi, 2023). With a definition, implicit bias means an optimizer regularized one term besides the training loss during the training, and we call this term the implicit bias of this optimizer. From our theory, we see any pointwise generalization bound results in a regularization (and thus an implicit bias) for algorithm design. For a bound in the form (3), considering the regularized ERM: $\hat{f} = \operatorname{argmin}_f \{ \mathbb{P}_n \ell(f;z) + C \sqrt{\frac{d(f) + \log(1/\delta)}{n}} \}$, its excess risk $\mathbb{P}\ell(\hat{f};z) - \mathbb{P}\ell(f^*;z)$ is bounded by (compared with any benchmark $f^* \in \mathcal{W}$),

$$\inf_{f \in \mathcal{F}} \left\{ \mathbb{P}\ell(f;z) + 2C\sqrt{\frac{d(f) + \log(1/\delta)}{n}} \right\} - \mathbb{P}\ell(f^*;z) \le 2C\sqrt{\frac{d(f^*) + \log(1/\delta)}{n}};$$

refer to Appendix E.2. This gives a problem-dependent upper bound $\sqrt{d(f^*)}/n$ (adapts to the optimal hypothesis f^*). Sparse linear models and matrix factorization explicitly impose sparsity/low-rank assumptions on f^* ; without strong convexity, the resulting statistical rates are unimprovable.

Table 1: Final-epoch Metrics of FCNs on MNIST. Supplementary explanations of columns: 1) Width -2^* means $h=2^*$; 2) Train Error; 3) Generalization error is defined as test error minus train error; 4) The spectral norm is the spectrally normalized margin bound of (Bartlett et al., 2017). It is the tightest norm—based bound in the literature to our knowledge; 5) Parameter Counts of the network; 6) VC dimension. We adopt a nearly tight VC-dimension bound from (Bartlett et al., 2019) and report $PL \log P$ for brevity (see Section 4.3); 7) R-D means our Riemannian Dimension.

Model	Train	Gen	Spectral Norm	# Parameters	VC dimension	R-D
Width-2 ⁶	0.0002	0.0205	3.146×10^{15}	5.961×10^{6}	9.299×10^{8}	6.433×10^{7}
Width- 2^7	0.0002	0.0187	2.695×10^{15}	6.167×10^{6}	9.641×10^{8}	6.097×10^{7}
Width-28	0.0000	0.0191	2.093×10^{15}	6.726×10^{6}	1.057×10^{9}	5.589×10^{7}
Width-29	0.0000	0.0186	2.401×10^{15}	8.434×10^{6}	1.345×10^{9}	5.316×10^{7}
Width- 2^{10}	0.0000	0.0215	4.816×10^{15}	1.421×10^{7}	2.340×10^{9}	5.266×10^{7}
Width- 2^{11}	0.0000	0.0160	1.001×10^{16}	3.520×10^{7}	6.116×10^{9}	4.972×10^{7}
Width- 2^{12}	0.0000	0.0210	1.466×10^{16}	1.149×10^{8}	2.133×10^{10}	4.803×10^{7}

Table 2: Final-Epoch Metrics of ResNets on CIFAR-10

Model	Train Error	Gen Error	# Parameters	VC dimension	R-D
ResNet-20	0.0016	0.0752	2.690×10^{5}	6.727×10^{7}	8.801×10^{6}
ResNet-32	0.0003	0.0695	4.630×10^{5}	1.933×10^{8}	9.992×10^{6}
ResNet-44	0.0001	0.0627	6.570×10^{5}	3.872×10^{8}	6.339×10^{6}
ResNet-56	0.0000	0.0637	8.510×10^{5}	6.507×10^{8}	5.200×10^{6}
ResNet-74	0.0000	0.0615	1.142×10^{6}	1.179×10^{9}	3.237×10^{6}
ResNet-110	0.0000	0.0576	1.724×10^{6}	2.723×10^{9}	2.583×10^{6}

4.3 COMPARISON WITH NORM BOUND, VC, AND NTK

Norm bound: Invoking the elementary bound $\log x \le \log(1+x) \le x$, we prove that the Riemannian–Dimension bound in Theorem 4 is *exponentially* tighter than the spectrally normalized bound in (Bartlett et al., 2017; Neyshabur et al., 2018); refer to Appendix E.3.

VC dimension: Let P be the number of weights and L be the number of layers, Bartlett et al. (2019) prove a nearly tight VC-dimension bound VCdim $\leq O\left(PL\log P\right)$, supported by a lower bound VCdim $\geq \Omega\left(PL\log(P/L)\right)$. This VC dimension bound is roughly equivalent to be $L\sum_{l=1}^{L}d_{l}d_{l-1}$. Our Riemannian-Dimension bound, by comparison, significantly sharpens this VC bound by removing the length L and replacing the width to be the effective rank.

NTK: Our approach uses exact non-perturbative expansion, which preserves the finite-scale geometry of deep networks, beyond NTK's Taylor linearizations that remains only valid in the infinitesimal-scale neighborhood around initialization. In effect, our method provides a finite-scale pointwise NTK in practical regimes, whereas NTK is a infinitesimal-scale theory resisting to the generalization theory.

5 EXPERIMENTS AT A GLANCE AND CONCLUSIONS

Tables 1 and 2 show that our *Riemannian dimension* is *orders of magnitude* tighter than prior complexity measures and—more strikingly—typically *decreases* as the number of parameters grows. Additional results in Appendix A further show that for FCNs on MNIST the effective feature rank can compress by as much as $300 \times$ (to 1/300 of width), and for ResNet on CIFAR-10 the *majority* of layers have (near) zero effective rank, explaining the benign effect of overparameterization. More details, tables, and discussion are deferred to Appendix A.

Collectively, our results establish a coherent, geometry-aware foundation for tractable and predictive generalization in fully connected neural networks. Addressing this core challenge required several technical innovations, and our experiments validate the resulting theoretical understanding.

REFERENCES

- Pierre Alquier et al. User-friendly introduction to pac-bayes bounds. *Foundations and Trends*® *in Machine Learning*, 17(2):174–303, 2024.
- Steve Awodey. Category theory, volume 52. OUP Oxford, 2010.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. *Advances in neural information processing systems*, 30, 2017.
- Peter L. Bartlett, Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *Journal of Machine Learning Research*, 20(63):1–17, 2019.
- Thomas Bendokat, Ralf Zimmermann, and P-A Absil. A grassmann manifold handbook: Basic geometry and computational aspects. *Advances in Computational Mathematics*, 50(1):6, 2024.
- Olivier Catoni. A pac-bayesian approach to adaptive classification. preprint, 840(2):6, 2003.
- Karel Devriendt, Hannah Friedman, Bernhard Reinke, and Bernd Sturmfels. The two lives of the grassmannian. *arXiv preprint arXiv:2401.03684*, 2024.
- Ilya Dumer, Mark S Pinsker, and Vyacheslav V Prelov. On coverings of ellipsoids in euclidean spaces. *IEEE transactions on information theory*, 50(10):2348–2356, 2004.
- Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv* preprint arXiv:1703.11008, 2017.
- Mathieu Even and Laurent Massoulié. Concentration of non-isotropic random tensors with applications to learning and empirical risk minimization. In *Conference on Learning Theory*, pp. 1847–1886. PMLR, 2021.
- Kenneth J Falconer. Techniques in fractal geometry, volume 3. Wiley Chichester, 1997.
- X Fernique. Regularite des trajectoires des fonctions aleatoires gaussiennes. *Ecole d'Eté de Probabilités de Saint-Flour IV—1974*, pp. 1–96, 1975.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *Information and Inference: A Journal of the IMA*, 9(2):473–504, 2020.
- Gene H Golub and Charles F Van Loan. Matrix computations. JHU press, 2013.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Geoffrey E Hinton and Drew Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pp. 5–13, 1993.
- Minyoung Huh, Hossein Mobahi, Richard Zhang, Brian Cheung, Pulkit Agrawal, and Phillip Isola. The low-rank simplicity bias in deep networks. *Transactions on Machine Learning Research*, 2021.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
 - Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

- Sanae Lotfi, Marc Finzi, Sanyam Kapoor, Andres Potapczynski, Micah Goldblum, and Andrew G
 Wilson. Pac-bayes compression bounds so tight that they can explain generalization. *Advances in Neural Information Processing Systems*, 35:31459–31473, 2022.
- Neil Lutz. A note on pointwise dimensions. arXiv preprint arXiv:1612.05849, 2016.
 - Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pp. 195–248. Springer, 1998.
- Vitali D Milman and Gideon Schechtman. Asymptotic theory of finite dimensional normed spaces:

 Isoperimetric inequalities in riemannian manifolds, volume 1200. Springer Science & Business Media, 1986.
 - Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representations*, 2018.
 - Christopher C Paige and Musheng Wei. History and generality of the cs decomposition. *Linear Algebra and its Applications*, 208:303–326, 1994.
 - Alain Pajor. Metric entropy of the grassmann manifold. *Convex Geometric Analysis*, 34(181-188): 0942–46013, 1998.
 - Liam Parker, Emre Onal, Anton Stengel, and Jake Intrater. Neural collapse in the intermediate hidden layers of classification neural networks. *arXiv preprint arXiv:2308.02760*, 2023.
 - Gilles Pisier. *The volume of convex bodies and Banach space geometry*, volume 94. Cambridge University Press, 1999.
 - Alessandro Rinaldo and Enxu Yan. Lecture 14: Uniform Bound via Rademacher Complexity. Scribed notes for CMU 36-755: Advanced Statistical Theory I, October 2016. URL https://www.stat.cmu.edu/~arinaldo/Teaching/36755/F16/Scribed_Lectures/AST_Oct19_Scribe.pdf.
 - Zhouxing Shi, Yihan Wang, Huan Zhang, J Zico Kolter, and Cho-Jui Hsieh. Efficiently computing local lipschitz constants of neural networks via bound propagation. *Advances in Neural Information Processing Systems*, 35:2350–2364, 2022.
 - Stanislaw J Szarek. Metric entropy of homogeneous spaces. arXiv preprint math/9701213, 1997.
 - Michel Talagrand. Regularity of gaussian processes. Acta Mathematica, 159:99–149, 1987.
 - Michel Talagrand. *The generic chaining: upper and lower bounds of stochastic processes*. Springer Science & Business Media, 2005.
 - Ramon Van Handel. Probability in high dimension. *Lecture Notes (Princeton University)*, 2(3):2–3, 2014.
 - Gal Vardi. On the implicit bias in deep-learning algorithms. *Communications of the ACM*, 66(6): 86–93, 2023.
- Roman Vershynin. High-dimensional probability, 2009.
- Peng Wang, Xiao Li, Can Yaras, Zhihui Zhu, Laura Balzano, Wei Hu, and Qing Qu. Understanding deep representation learning via layerwise feature compression and discrimination. *arXiv* preprint *arXiv*:2311.02960, 2025.
- Wikipedia contributors. Change of variables. https://en.wikipedia.org/wiki/ Change_of_variables, a.
 - Wikipedia contributors. Min-max theorem. https://en.wikipedia.org/wiki/Min-max_theorem, b.
- Wikipedia contributors. Square root of a matrix. https://https://en.wikipedia.org/wiki/Square_root_of_a_matrix, c.

Wikipedia contributors. Volume of an n-ball. https://en.wikipedia.org/wiki/ Volume_of_an_n-ball, d. David P Woodruff et al. Sketching as a tool for numerical linear algebra. Foundations and Trends® in Theoretical Computer Science, 10(1–2):1–157, 2014. Yunbei Xu and Assaf Zeevi. Towards optimal problem dependent generalization error bounds in statistical learning theory. *Mathematics of Operations Research*, 50(1):40–67, 2025.

Table 3: Final-epoch Metrics of FCNs on MNIST. Supplementary explanations of columns: 1) Width -2^* means $h=2^*$; 2) Train Error; 3) Generalization error is defined as test error minus train error; 4) The spectral norm is the spectrally normalized margin bound of (Bartlett et al., 2017). It is the tightest norm-based bound in the literature to our knowledge; 5) Parameter Counts of the network; 6) VC dimension. We adopt a nearly tight VC-dimension bound from (Bartlett et al., 2019) and report $PL \log P$ for brevity (see Section 4.3); 7) R-D means our Riemannian Dimension.

Model	Train	Gen	Spectral Norm	# Parameters	VC dimension	R-D
Width-26	0.0002	0.0205	3.146×10^{15}	5.961×10^{6}	9.299×10^{8}	6.433×10^{7}
Width- 2^7	0.0002	0.0187	2.695×10^{15}	6.167×10^{6}	9.641×10^{8}	6.097×10^{7}
Width-28	0.0000	0.0191	2.093×10^{15}	6.726×10^{6}	1.057×10^{9}	5.589×10^{7}
Width-29	0.0000	0.0186	2.401×10^{15}	8.434×10^{6}	1.345×10^{9}	5.316×10^{7}
Width- 2^{10}	0.0000	0.0215	4.816×10^{15}	1.421×10^{7}	2.340×10^{9}	5.266×10^{7}
Width- 2^{11}	0.0000	0.0160	1.001×10^{16}	3.520×10^{7}	6.116×10^{9}	4.972×10^{7}
Width- 2^{12}	0.0000	0.0210	1.466×10^{16}	1.149×10^8	2.133×10^{10}	4.803×10^{7}

Table 4: Final-Epoch Metrics of ResNets on CIFAR-10

Model	Train Error	Gen Error	# Parameters	VC dimension	R-D
ResNet-20	0.0016	0.0752	2.690×10^{5}	6.727×10^{7}	8.801×10^{6}
ResNet-32	0.0003	0.0695	4.630×10^{5}	1.933×10^{8}	9.992×10^{6}
ResNet-44	0.0001	0.0627	6.570×10^{5}	3.872×10^{8}	6.339×10^{6}
ResNet-56	0.0000	0.0637	8.510×10^{5}	6.507×10^{8}	5.200×10^{6}
ResNet-74	0.0000	0.0615	1.142×10^{6}	1.179×10^{9}	3.237×10^{6}
ResNet-110	0.0000	0.0576	1.724×10^{6}	2.723×10^{9}	2.583×10^{6}

A EXPERIMENTS

We evaluate our Riemannian Dimension on two standard architectures—Fully Connected Networks (FCNs) and ResNets, using two benchmark datasets—MNIST (LeCun et al., 1998) and CIFAR-10 (Krizhevsky, 2009), respectively. We consider a 9-hidden-layer FCN architecture, where, except for the fixed layers, hidden layers share a common width h, with $h \in \{2^6, 2^7, 2^8, 2^9, 2^{10}, 2^{11}, 2^{12}\}$. Increasing h monotonically enlarges both layer widths and model sizes. We adopt canonical ResNet architectures—ResNet-20, ResNet-32, ResNet-44, ResNet-56, ResNet-74 and ResNet-110—which differ only in the number of residual blocks per stage while maintaining the same overall architecture (three-stage, basic-block design) as introduced by (He et al., 2016). These ResNet architectures provides a clean capacity sweep via depth. In what follows, we organize experiments around the two complementary regimes—width scaling on FCNs and depth scaling on ResNets.

This design lets us systematically study three central questions in modern deep learning: (i) why does overparameterization often improve generalization? (ii) how does feature learning evolve during training? and (iii) what implicit regularization is encoded by the baseline optimizer? Detailed experimental setups are deferred to to the end of this section.

A.1 RIEMANNIAN DIMENSION EXPLAINS OVERPARAMETERIZATION

This section studies why does overparameterization—despite exploding model capacity—often improve generalization. We investigate this paradox by tracking our Riemannian Dimension across models with varying parameter counts, asking whether more parameters truly enlarge capacity or instead deduce complexity.

Final-epoch metrics of FCNs on MNIST and ResNets on CIFAR-10 are reported in Table 3 and Table 4, respectively. In these Tables, the train error quickly collapses to zero for sufficiently large models, confirming their expressive capacity. Consistently, the generalization can continue to be improved as parameters increase, especially on ResNets (Table 4). This phenomenon means the overfitting does not appear and reflects a paradoxical truth of deep learning: over-parameterization

Table 5: Final-epoch Effective Ranks for FCNs on MNIST, where Width -2^* means $h=2^*$, and where for the form A/B, A represents the effective rank and B represents the original dimension, and where Layer-1 means the input layer.

Metric	Width-26	Width-2 ⁷	Width-28	Width-29	Width-2 ¹⁰	Width-2 ¹¹	Width-2 ¹²
Layer-1	713/763	712/763	710/763	710/763	707/763	707/763	704/763
Layer-2	2048/2048	2044/2048	2042/2048	2048/2048	2047/2048	2048/2048	2048/2048
Layer-3	2048/2048	2045/2048	2037/2048	2019/2048	1925/2048	1460/2048	1009/2048
Layer-4	61/64	97/128	92/256	85/512	79/1024	79/2048	59/4096
Layer-5	23/64	43/128	34/256	33/512	28/1024	26/2048	22/4096
Layer-6	20/64	24/128	20/256	21/512	19/1024	18/2048	15/4096
Layer-7	15/64	18/128	17/256	15/512	15/1024	14/2048	13/4096
Layer-8	15/64	14/128	15/256	11/512	13/1024	13/2048	12/4096
Layer-9	14/64	14/128	15/256	13/512	13/1024	12/2048	12/4096
Layer-10	13/64	13/128	12/256	14/512	12/1024	13/2048	14/4096
Total	4970	5024	4994	4969	4858	4390	3908

Table 6: Final-epoch Effective Ranks for ResNets on CIFAR-10, where for the form A/B, A represents the effective rank and B represents the original dimension, and where Layer-0% means the input layer.

Metric	ResNet-20	ResNet-32	ResNet-44	ResNet-56	ResNet-74	ResNet-110
Layer-0%	384/3072	384/3072	17/3072	0/3072	0/3072	0/3072
Layer- 25%	2048/16384	2048/16384	7/16384	1/16384	0/16384	0/16384
Layer-50%	1024/8192	1024/8192	1024/8192	227/8192	0/8192	0/8192
Layer-75%	512/4096	512/4096	512/4096	512/4096	58/4096	0/4096
Layer-100%	8/64	8/64	8/64	8/64	8/64	8/64
Total	23432	37768	27564	16294	11401	6925

is not a curse, but can benefit the generalization. However, classical complexity measures—e.g., the spectral norm and the VC dimension, often scale exponentially as the parameter count grows. Notably, the spectral norm is about 10^6 times larger than the VC dimension and seems to be a worse complexity measure (see Table 3). The two measures therefore struggle to explain the generalization of modern overparameterized networks. In contrast, our Riemannian Dimension exhibits a consistent downward trend as model size grows—both under width scaling (last column of Table 3) and depth scaling (last column of Table 4), and it is about 10^3 times smaller than the VC dimension, suggesting that the effective dimension—not raw parameter count—is the true indicator of generalization in deep learning. In summary, increased parameterization is associated with reduced effective model complexity, and Riemannian Dimension faithfully characterizes this phenomenon.

A.2 FEATURE LEARNING COMPRESSES EFFECTIVE RANK

We investigate the dynamics of feature learning by monitoring the effective rank of the feature Gram matrices $F_{l-1}F_{l-1}^{\top}$, with the normalization $|L||W||_F^2\Pi_{i>l}\|W_i\|_{\text{op}}^2$ dictated by our theory. Empirical results appear in Tables 5, 6 and Figure 1.

Experimental results reveal some clear patterns: (1) As training proceeds, the effective ranks of feature grams decreases sharply after a short transient; refer to Figure 1. (2) Increased parameter counts, both under width scaling (FCNs) and depth scaling (ResNets), foster compressing effective ranks of feature grams in both the rate and the degree; refer to Figure 1. (3) On the largest FCN, the degree of effective rank compression can reach as much as 1/300, which explains why the Riemannian Dimension can achieve such a significant improvement over the VC dimension; refer to Table 5. While on the largest ResNet, the effective ranks of the vast majority of layers compress to zero, which explains why deeper networks can, paradoxically, exhibit a smaller Riemannian Dimension; refer to Table 6. These experimental results indicates that feature learning steadily reduces

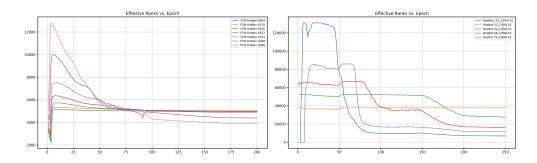


Figure 1: Effective Rank evolutions of FCNs on MNIST (left) and ResNets on CIFAR-10 (right) across the training

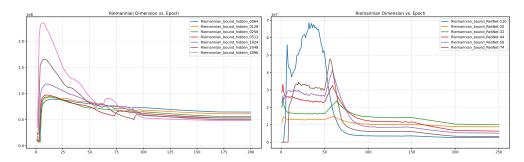


Figure 2: Riemannian Dimension evolutions of FCNs on MNIST (left) and ResNets on CIFAR-10 (right) across the training

the intrinsic dimensionality of features over training and aim to learn a lower-dimensional feature manifold, and the overparameterization intensifies this reduction.

A.3 SGD FINDS LOW RIEMANNIAN DIMENSION POINT

Related literature has shown that various norms are implicit bias of optimizers, but typically limited to linear models (Vardi, 2023). This section studies whether SGD with momentum, in modern deep learning, implicit regularized Riemannian Dimension across training dynamics. We examine whether this optimizer preferentially converge to solutions with lower Riemannian Dimension point, and the experimental results are presented in Figure 2.

Empirical results show a repeatable pattern across the architectures: SGD with momentum drives the networks toward solutions with lower intrinsic Riemannian Dimension complexity, after an early transient; refer to Figure 2. Notably, Riemannian Dimension drops by orders of magnitude, whereas VC dimension remains essentially unchanged. The alignment between optimization dynamics and complexity control supports the view that SGD with momentum implicitly regularizes the Riemannian Dimension. Therefore, optimization is not merely as a mechanism for convergence; it is a primary driver of generalization through its systematic preference for low-complexity solutions. Riemannian Dimension provides a practical and theoretically grounded lens through which the implicit bias of optimizes in machine learning can be quantitatively assessed.

Experimental Setup. We introduce detailed experimental setups. We evaluate our Riemannian Dimension bound on two standard architectures—Fully Connected Networks (FCNs) and ResNets, using two benchmark datasets—MNIST (LeCun et al., 1998) and CIFAR-10 (Krizhevsky, 2009), respectively. The architecture of FCNs: we consider a 9-hidden-layer FCN in which the first two hidden layers have width 2^{11} and the remaining seven hidden layers share a common width h, with $h \in \{2^6, 2^7, 2^8, 2^9, 2^{10}, 2^{11}, 2^{12}\}$. The output layer is a linear classifier mapping to 10 logits, and we use ReLU as the activation and use PyTorch's default initialization (Kaiming uniform for ReLU). Increasing h monotonically enlarges both layer widths and the total parameter count, yielding a clean capacity sweep at fixed depth. The architecture of ResNets: we adopt the canonical ResNet

architectures, ResNet-20, ResNet-32, ResNet-44, ResNet-56, ResNet-74, and ResNet-110, which differ only in the number of residual blocks per stage while maintaining the same overall architecture (three-stage, basic-block design) as introduced by (He et al., 2016). Following the practice of (He et al., 2016), we apply BatchNorm and ReLU after each convolution, with shortcut connections added as needed, and a global average pooling layer precedes the final linear classifier. These ResNet architectures provides a clean capacity sweep via depth.

We adopt standard training pipelines widely used in the benchmarks. (1) The training Protocol of FCNs is: SGD with momentum optimizer where momentum = 0.9, learning rate = 0.01, and weight decay = 5×10^{-4} ; 200 epochs and 128 batch size; a step decay at epochs $\{100, 170\}$, where the learning rate is scaled by $\times 0.1$. (2) The training Protocol of ResNets is: SGD with momentum optimizer where momentum = 0.9, learning rate = 0.1, and weight decay = 5×10^{-4} ; 250 epochs and 128 batch size; a step decay at epochs $\{50, 150, 200\}$, where the learning rate is scaled by $\times 0.1$; Following practical training conditions, we apply standard data augmentation on CIFAR-10: random horizontal flips and 4-pixel random crops with zero-padding.

In the experiments of FCNs and ResNets, to enable layerwise analysis of the evolving feature representations and support our computation of Riemannian Dimension, we register forward hooks on all nonlinearity layers. For layers followed by pooling, we replace the last recorded ReLU activation with the corresponding pooled output. We also pre-register the input hook to capture the feature matrix of the data. These hooks ensures precise extraction of nonlinearity activations at each depth throughout training. We set the hyper-parameter ε via a one–dimensional ternary-search procedure: at the end of each training stage we perform a 500-step ternary search for FCNs and a 50-step ternary search for ResNets over the admissible interval

 $[\sqrt{1/n}, \max_{l=1,\dots,L} \sqrt{\frac{2L\lambda_{\max}(F_{l-1}F_{l-1}^{\mathsf{T}})\cdot||W||_F^2\Pi_{i>l}\|W_i\|_{op}^2}{n}}]$. The search selects the value that minimizes our one-shot Riemannian Dimension-based generalization error bound (Theorem 1). We note that tighter bounds could be achieved with more refined optimization procedures on ε . For FCNs, we compute full feature gram matrices. While for ResNets, the feature matrix F is formed by flattening each activation map into a vector of dimension $d=C\cdot H\cdot W$, where C,H,W are the channel, height, and width of the feature map respectively. To align with our theory, we simplify ResNets to fully connected (feed-forward) networks when computing our bound; we apply the same simplification to the associated VC-dimension and parameter-count calculations to maintain consistency. To avoid out-of-memory in computing full feature gram matrices in high-dimensional convolutional layers, we use the standard Gaussian sketching approximation, where each feature gram matrix uses a Gaussian sketch with parameter $r=\min(8192,\lfloor d/8 \rfloor)$ (Woodruff et al., 2014). By standard subspace-embedding guarantees, such Gaussian sketches preserve Gram quadratic forms—and hence the spectra—of the feature matrices with high probability, introducing only negligible distortion and leaving our conclusions unchanged (Woodruff et al., 2014).

B PROOFS FOR THE NATURE OF GENERALIZATION (SECTION 2)

B.1 PROOF OF THEOREM 1 (ONE-SHOT GENERALIZATION BOUND)

To prove our pointwise-dimension generalization bound, we start with a novel combination of symmetrization and conditional PAC-Bayes, which enables data-dependent prior and empirical L_1/L_2 metric control. In the next step, we show how a standard analysis of symmetrized loss extends this result to original losses, incurring at most a constant factor of 2 in the bound.

B.1.1 SYMMETRIZED LOSSES AND CONDITIONAL PAC-BAYES

As discussed above, we start with a one-shot approximation bound for symmetrized loss. Now define $g(f;z,\xi)=\xi\ell(f;z)$, where ξ is an independent Rademacher variable, a symmetric binary random variable taking values with $\Pr(\xi=+1)=\Pr(\xi=-1)=\frac{1}{2}$. Denote the population and empirical distributions of ξ as $\mathbb Q$ and $\mathbb Q_n$. By independence assumption we know that Q is also the conditional distribution of (z,ξ) conditioned on z, thus we can use the product distribution (such as $\mathbb Q_n\otimes\mathbb P_n$) to denote the joint distributions of (z,ξ) . A key property of symmetrized loss is that it

preserves the empirical L_2 and L_1 metric of the original loss:

$$\varrho_{n,\ell}(f,f') := \sqrt{\mathbb{P}_n[(\ell(f;z) - \ell(f';z))^2]} = \sqrt{\mathbb{Q}_n \otimes \mathbb{P}_n[(g(f;z,\xi) - g(f';z,\xi))^2]},$$
(11)

$$L_1(\mathbb{P}_n)(f,f') = \mathbb{P}_n|\ell(f;z) - \ell(f';z)| = \mathbb{Q}_n \otimes \mathbb{P}_n|\ell(f;z) - \ell(f';z)|. \tag{12}$$

because for Rdemacher variables ξ , $(\ell(f;z) - \ell(f';z))^2 = [\xi(\ell(f;z) - \ell(f';z))]^2 = (g(f;z,\xi) - g(f';z,\xi))^2$, and $|\ell(f;z) - \ell(f';z)| = |\xi(\ell(f;z) - \ell(f';z))| = |g(f;z,\xi) - g(f';z,\xi)|$.

We are interested in the conditional generalization gap (conditioned on every fixed realization of $\{z_i\}_{i=1}^n$ regardless of \mathbb{P}):

$$\mathbb{Q} \otimes \mathbb{P}_n g(f; z, \xi) - \mathbb{Q}_n \otimes \mathbb{P}_n g(f; z, \xi) = 0 - \frac{1}{n} \sum_{i=1}^n g(f; z_i, \xi_i).$$

For technical reasons (the later application of McDiarmid's inequality via bounded difference to prove Theorem 1), we work with $L_1(\mathbb{P}_n)$ metric now (which is a smaller metric than $L_2(\mathbb{P}_n)$ so the bound is more general).

Lemma 4 Let π be any data-dependent prior on a hypothesis class \mathcal{F} (π can depend on z_1, \dots, z_n), and let $\ell \colon \mathcal{F} \times \mathcal{Z} \to [0,1]$ be a bounded loss. Fix confidence $\delta \in (0,1)$ and sample size n. Then for all $\varepsilon > 0$, with probability at least $1 - \delta$ over n i.i.d. draws ξ_1, \dots, ξ_n from \mathbb{Q} , uniformly over all $f \in \mathcal{F}$,

$$0 - \frac{1}{n} g(f; z_i, \xi_i) \leq \varepsilon + \sqrt{\frac{\log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f, \varepsilon))} + \log \frac{1}{\delta}}{2n}}.$$

Proof of Lemma 4: the proof consists of three steps: 1. Conditional PAC-Bayes; 2. Approximating the Deterministic Hypothesis; and 3. Taking Uniform Posterior in ε -Ball.

1. Conditional PAC-Bayes. Applying Catoni's PAC-Bayes uniform bound (Catoni, 2003) (Lemma 9 in our Auxiliary Lemma Part) to the population and empirical distribution of ξ , we have that with probability at least $1-\delta$, uniformly over every $\mu\in\Delta(\mathcal{F})$, we have the "random hypothesis" bound

$$\langle \mu, 0 - \mathbb{Q}_n \otimes \mathbb{P}_n g(\cdot; z, \xi) \rangle \le \sqrt{\frac{\mathrm{KL}(\mu, \pi) + \log \frac{1}{\delta}}{2n}},$$
 (13)

where the absolute constant comes from the loss $g \in [-1,1]$. The D_{KL} divergence is defined by $\mathrm{KL}(\mu \| \pi) = \int_{\mathcal{F}} \log \left(\frac{d\mu}{d\pi}(f) \right) \mu(df)$.

2. Approximating the Deterministic Hypothesis. So on the event that the above conditional PAC-Bayes inequality (13) holds, with probability at least $1-\delta$, we have that uniformly over every random $\mu \in \Delta(\mathcal{F})$ and at the same time uniformly over every deterministic $f \in \mathcal{F}$, for every $\eta > 0$, we have the "deterministic hypothesis" bound: for any $\eta > 0$,

$$0 - \mathbb{Q}_{n} \otimes \mathbb{P}_{n}g(f; z, \xi)$$

$$= \langle \mu, 0 - \mathbb{Q}_{n} \otimes \mathbb{P}_{n}g(\cdot; z, \xi) \rangle + \langle \mu, \mathbb{Q}_{n} \otimes \mathbb{P}_{n}[g(\cdot; z, \xi) - g(f; z, \xi)] \rangle$$

$$\leq \frac{\eta}{8} + \frac{\text{KL}(\mu, \pi) + \log \frac{1}{\delta}}{\eta n} + \langle \mu, \mathbb{Q}_{n} \otimes \mathbb{P}_{n}[|g(\cdot; z, \xi) - g(f; z, \xi)|] \rangle$$

$$= \frac{\eta}{8} + \frac{\text{KL}(\mu, \pi) + \log \frac{1}{\delta}}{\eta n} + \langle \mu, L_{1}(\mathbb{P}_{n})(\cdot, f) \rangle$$
(14)

where the first identity holds because $\langle \mu, \mathbb{Q}_n \otimes \mathbb{P}_n g(f; z, \xi) \rangle = g(f; z, \xi)$; the inequality uses conditional PAC-Bayes (13) to bound the first term and use absolute values to bound the second term; and the last identity uses metric preservation of symmetrization (identity (11)). On the event that (13) holds, the new inequality (14) holds simultaneously not only for every random μ (due to PAC-Bayes) but also for every deterministic f, because on such event, the new inequality (14) does not use any new probabilistic inequality besides (13).

Taking Uniform Posterior in ε **-ball.** Given any prior π on \mathcal{F} , for any $f \in \mathcal{F}$, we take μ to be the uniform distribution over $B_{L_1(\mathbb{P}_n)}(f,\varepsilon)$:

$$\mu = \text{Unif}(B_{L_1(\mathbb{P}_n)}(f, \varepsilon)). \tag{15}$$

We emphasize that using this uniform measure over the ε -ball is "essentially optimal" in that it reproduces the same analytical upper bound—namely the conclusion of Lemma 4—as would be obtained by employing the Gibbs distribution that minimizes the bound in (14).

This "essentially optimal" choice (15) of μ yields

$$\frac{d\mu}{d\pi}(f) = \begin{cases} \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f,\varepsilon))}, & f \in B_{L_1(\mathbb{P}_n)}(f,\varepsilon), \\ 0, & f \notin B_{L_1(\mathbb{P}_n)}(f,\varepsilon). \end{cases}$$

the KL divergence is

$$KL(\mu \| \pi) = \int_{\mathcal{F}} \log \left(\frac{d\mu}{d\pi}(f) \right) \mu(df)$$

$$= \int_{B_{L_1(\mathbb{P}_n)}(f,\varepsilon)} \log \left(\frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f,\varepsilon))} \right) \mu(df)$$

$$= \log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f,\varepsilon))} \underbrace{\int_{B_{L_1(\mathbb{P}_n)}(f,\varepsilon)} \mu(df)}_{=1}$$

$$= \log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f,\varepsilon))}.$$
(16)

Because (14) holds simultaneously for every w and every $\mu = \text{Unif}(B_{L_1(\mathbb{P}_n)}(f,\varepsilon))$, we can apply (14) and substitute KL equality (16) back to (14). Minimizing over η , this gives:

$$0 - \mathbb{Q}_n g(f; z, \xi) \le \sqrt{\frac{\log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f, \varepsilon))} + \log \frac{1}{\delta}}{2n}} + \varepsilon.$$

The proof develops the generalization error as an optimization problem on bias and variance. Through introducing symmetrized losses and a conditioning technique, we transformer the biasvariance optimization to a PAC-Bayes optimization problem on the posterior and prior. We solve it optimally by establishing matching upper and lower bounds; refer to the following lemma.

Lemma 5 (Optimality of the ε -ball uniform posterior for the PAC–Bayes optimization) Define the PAC–Bayes optimization functional

$$V_f(\eta; P_n, \pi) := \inf_{\mu \ll \pi} \left\{ \frac{\eta}{8} + \frac{\mathrm{KL}(\mu, \pi) + \log \frac{1}{\delta}}{\eta n} + \langle \mu, L_1(\mathbb{P}_n)(\cdot, f) \rangle \right\}.$$

Then for every $f \in \mathcal{F}$, $\eta > 0$, and $\varepsilon > 0$:

$$V_f(\eta; P_n, \pi) \ge \frac{\eta}{8} + \frac{\log \frac{1}{\delta}}{\eta n} + \min \left\{ \frac{1}{\eta n} \log \frac{1}{\pi (B_{L_1(\mathbb{P}_n)}(f, \varepsilon))}, \varepsilon \right\} - \frac{\log 2}{\eta n}. \tag{17}$$

Consequently, for every $f \in \mathcal{F}$, $\eta > 0$, and $\varepsilon > 0$,

$$\frac{\eta}{8} + \frac{\log \frac{1}{\delta}}{\eta n} + \min \left\{ \frac{\log \frac{1}{\pi \left(B_{L_1(\mathbb{P}_n)}(f,\varepsilon)\right)}}{\eta n}, \varepsilon \right\} - \frac{\log 2}{\eta n} \le V_f(\eta; P_n, \pi) \le \frac{\eta}{8} + \frac{\log \frac{1}{\pi \left(B_{L_1(\mathbb{P}_n)}(f,\varepsilon)\right)} + \log \frac{1}{\delta}}{\eta n} + \varepsilon.$$
(18)

Proof of Lemma 5 The Donsker-Varadhan variational identity states that for any measurable h,

$$-\log \int e^h d\pi = \inf_{\mu \ll \pi} \Big\{ D_{\mathrm{KL}}(\mu \| \pi) - \int h d\mu \Big\}.$$

Apply it with $h = -\eta n L_1(\mathbb{P}_n)(\cdot, f)$ to obtain

$$-\log \int e^{-\eta n L_1(\mathbb{P}_n)(\cdot, f)} d\pi = \inf_{\mu \ll \pi} \Big\{ D_{\mathrm{KL}}(\mu \| \pi) + \int \eta n L_1(\mathbb{P}_n)(\cdot, f) d\mu \Big\},\,$$

which implies that

$$\frac{\eta}{8} + \frac{\log \frac{1}{\delta}}{\eta n} - \frac{1}{\eta n} \log \int e^{-\eta n L_1(\mathbb{P}_n)(\cdot, f)} d\pi = \inf_{\mu \ll \pi} \left\{ \frac{\eta}{2} + \frac{\mathrm{KL}(\mu, \pi) + \log \frac{1}{\delta}}{\eta n} + \langle \mu, L_1(\mathbb{P}_n)(\cdot, f) \rangle \right\},\tag{19}$$

By splitting the dual integral,

$$\int e^{-\eta n L_1(\mathbb{P}_n)(\cdot,f)} d\pi = \int_{B_{L_1(\mathbb{P}_n)}(f,\varepsilon)} e^{-\eta n L_1(\mathbb{P}_n)(\cdot,f)} d\pi + \int_{B_{L_1(\mathbb{P}_n)}(f,\varepsilon)^c} e^{-\eta n L_1(\mathbb{P}_n)(\cdot,f)} d\pi
\leq \pi \left(B_{L_1(\mathbb{P}_n)}(f,\varepsilon) \right) + e^{-\eta n \varepsilon} \left(1 - \pi \left(B_{L_1(\mathbb{P}_n)}(f,\varepsilon) \right) \right)
\leq \pi \left(B_{L_1(\mathbb{P}_n)}(f,\varepsilon) \right) + e^{-\eta n \varepsilon},$$

where we have used $e^{-\eta nL_1(\mathbb{P}_n)(\cdot,f)} \leq 1$ on $B_{L_1(\mathbb{P}_n)}(f,\varepsilon)$ and $e^{-\eta nL_1(\mathbb{P}_n)(\cdot,f)} \leq e^{-\eta n\varepsilon}$ on $B_{L_1(\mathbb{P}_n)}(f,\varepsilon)^c$. Hence

$$V_f(\eta; P_n, \pi) \ge \frac{\eta}{8} + \frac{\log \frac{1}{\delta}}{\eta n} - \frac{1}{\eta n} \log \left(\pi \left(B_{L_1(\mathbb{P}_n)}(f, \varepsilon) \right) + e^{-\eta n \varepsilon} \right). \tag{20}$$

The simplified form (17) follow from $a+b \le 2 \max\{a,b\}$ or equivalently $-\log(a+b) \ge -\log 2 + \min\{-\log a, -\log b\}$ on (20). Combining (14), (16) and (17) yields the sandwich (18).

B.1.2 ORIGINAL LOSSES

In the next step, via symmetrization, we extend the bound for symmetrized losses (Lemma 4 to original losses.

Proof of Theorem 1: The proof consists of three steps, following the standard symmetrization routine in controlling empirical processes: 1. Obtain In-Expectation Bound for Offset Symmetrized Losses; 2. Bounding Offset Symmetrized Losses by Offset Original Losses; 3. Applying McDiarmid's Inequality to Get High-Probability Bound.

1. Obtain In-Expectation Bound for Offset Symmetrized Losses. The bound in Lemma 4 implies

$$\Pr\left(\sup_{f\in\mathcal{F}}\left\{0-\frac{1}{n}\sum_{i=1}^{n}g(f;z_{i},\xi_{i})-\sqrt{\frac{\log\frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f,\varepsilon))}+\log\frac{1}{\delta}}{2n}}-\varepsilon\right\}>0\right)\leq\delta$$

which implies (by $\mathbb{Q}\otimes \mathbb{P}g(f;z,\xi)=0$ and $\sqrt{a+b}\leq \sqrt{a}+\sqrt{b}$)

$$\Pr\left(\sup_{f\in\mathcal{F}}\left\{(\mathbb{Q}\otimes\mathbb{P}-\mathbb{Q}_n\mathbb{P}_n)g(f;z,\xi)-\sqrt{\frac{\log\frac{1}{\pi(B_{L_1}(\mathbb{P}_n)(f,\varepsilon))}}{2n}}-\varepsilon\right\}>\sqrt{\frac{\log\frac{1}{\delta}}{n}}\right)\leq\delta.$$

Denoted by
$$\sup_{f \in \mathcal{F}} \left\{ (\mathbb{Q} \otimes \mathbb{P} - \mathbb{Q}_n \mathbb{P}_n) g(f; z, \xi) - \sqrt{\frac{\log \frac{1}{\pi(B_{L_1}(\mathbb{P}_n)^{(f, \varepsilon)})}}{2n}} - \varepsilon \right\} = Y$$
, then

$$\mathbb{E}Y \leq \mathbb{E}Y_+ = \int_0^\infty \Pr(Y > t) dt \leq \int_0^\infty e^{-nt^2} dt = \sqrt{\frac{\pi}{4n}} < \sqrt{\frac{1}{n}}.$$

Thus, we now get an expectation bound.

$$\mathbb{E}\sup_{f\in\mathcal{F}}\left\{(\mathbb{Q}\otimes\mathbb{P}-\mathbb{Q}_n\mathbb{P}_n)g(f;z,\xi)-\sqrt{\frac{\log\frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f,\varepsilon))}}{2n}}-\varepsilon-\sqrt{\frac{1}{n}}\right\}\leq 0. \tag{21}$$

2. Bounding Offset Symmetrized Losses by Offset Original Losses. Now we apply the following analysis, standard in using symmetrization to control empirical processes:

$$\begin{split} &\mathbb{E}_{z} \sup_{f \in \mathcal{F}} \left\{ (\mathbb{P} - \mathbb{P}_{n}) \ell(f; z) - 2 \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - 2\varepsilon - \sqrt{\frac{4}{n}} \right\} \\ &= \mathbb{E}_{z} \sup_{f \in \mathcal{F}} \mathbb{E}_{z'} \left\{ \mathbb{P}_{n} [(\ell(f; z') - \ell(f; z))] - 2 \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - 2\varepsilon - \sqrt{\frac{4}{n}} \right\} \\ &\leq \mathbb{E}_{z, z'} \sup_{f \in \mathcal{F}} \left\{ \mathbb{P}_{n} [(\ell(f; z') - \ell(f; z))] - 2 \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - 2\varepsilon - \sqrt{\frac{4}{n}} \right\} \\ &= \mathbb{E}_{z, z', \xi} \sup_{f \in \mathcal{F}} \left\{ \mathbb{Q}_{n} \otimes \mathbb{P}_{n} [\xi(\ell(f; z') - \ell(f; z))] - 2 \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - 2\varepsilon - \sqrt{\frac{4}{n}} \right\} \\ &\leq \mathbb{E}_{z', \xi} \sup_{f \in \mathcal{F}} \left\{ \mathbb{Q}_{n} \otimes \mathbb{P}_{n} \xi \ell(f; z') - \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - \varepsilon - \sqrt{\frac{1}{n}} \right\} \\ &+ \mathbb{E}_{z, \xi} \sup_{f \in \mathcal{F}} \left\{ \mathbb{Q}_{n} \otimes \mathbb{P}_{n} (-\xi \ell(f; z)) - \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - \varepsilon - \sqrt{\frac{1}{n}} \right\} \\ &= 2 \mathbb{E}_{z, \xi} \sup_{f \in \mathcal{F}} \left\{ (\mathbb{Q} \otimes \mathbb{P} - \mathbb{Q}_{n} \otimes \mathbb{P}_{n}) g(f; z, \xi) - \sqrt{\frac{\log \frac{1}{\pi(B_{L_{1}(\mathbb{P}_{n})}(f, \varepsilon))}}{2n}} - \varepsilon - \sqrt{\frac{1}{n}} \right\} \\ &\leq 0 \end{split}$$

where the first inequality is because sup is a convex function; the second equality is because $\ell(f;z') - \ell(f;z)$ has a symmetric distribution; the second inequality is because $\sup_x (h_1(x) + h_2(x)) \le \sup_x h_1(x) + \sup_x h_2(x)$; and the last equality is because $\xi \ell(f;z')$ has the same distribution with $\xi \ell(f,z)$; and the last inequality is by (21). We conclude that

$$\mathbb{E}_z \sup_{f \in \mathcal{F}} \left\{ (\mathbb{P} - \mathbb{P}_n) \ell(f; z) - \inf_{\varepsilon > 0} \left\{ \sqrt{\frac{2 \log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f, \varepsilon))}}{n}} + 2\varepsilon \right\} - \sqrt{\frac{4}{n}} \right\} \leq 0$$

3. Applying McDiarmid's Inequality to Get High-Probability Bound. Finally, we apply McDiarmid's inequality (Lemma 10) to get the high probability bound. By an arbitrary change of the sample $s = \{z_1, ..., z_n\}$ to $S' = \{z_1, ..., z_j', ..., z_n\}$, the bounded difference of $\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P})$

$$\mathbb{P}_n)\ell(f;z) - \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2\log\frac{1}{\pi(B_{L_1}(\mathbb{P}_n)(f,\varepsilon))}}{n}} + 2\varepsilon \right\} - \sqrt{\frac{4}{n}} \text{ is at most } 3/n. \text{ Specifically, define } 1 + 2\varepsilon$$

$$\psi(f;S) = (\mathbb{P} - \mathbb{P}_n)\ell(f;z) - \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2\log \frac{1}{\pi(B_{L_1}(\mathbb{P}_n)(f,\varepsilon))}}{n}} + 2\varepsilon \right\} - \sqrt{\frac{4}{n}},$$

we want to show

$$|\sup_{f} \psi(f; S) - \sup_{f} \psi(f; S')| \le \frac{3}{n}.$$

Clearly, to prove (22), it suffices to show that for every fixed $w \in \mathcal{F}$,

$$|\psi(f,S) - \psi(f,S')| \le \frac{3}{n}, \quad \forall \text{ fixed } f. \tag{22}$$

Denote $L_1(S)$ and $L_1(S')$ to be the empirical L_2 norms with S and S', respectively. We have that $|\psi(f;S)-\psi(f;S')|$ is bounded by

$$\left| \frac{\ell(f;z') - \ell(f;z)}{n} + \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2\log \frac{1}{\pi(B_{L_1(S')}(f,\varepsilon))}}{n}} + 2\varepsilon \right\} - \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2\log \frac{1}{\pi(B_{L_1(S)}(f,\varepsilon))}}{n}} + 2\varepsilon \right\} \right|, \tag{23}$$

For all $\varepsilon > 1/n$, as ℓ is bounded in [0, 1],

$$B_{L_1(S)}(f,\varepsilon-1/n) \subseteq B_{L_1(S')}(f,\varepsilon) \subseteq B_{L_1(S)}(f,\varepsilon+1/n),$$

so we have

$$\log \frac{1}{\pi(B_{L_1(S)}, \varepsilon + 1/n)} \le \log \frac{1}{\pi(B_{L_1(S')}, \varepsilon)} \le \log \frac{1}{\pi(B_{L_1(S)}, \varepsilon - 1/n)}$$

and thus (taking $\varepsilon_1 = \varepsilon - 1/n$ and $\varepsilon_2 = \varepsilon + 1/n$.

$$\begin{split} \inf_{\varepsilon_1 > 2/n} \left\{ \sqrt{\frac{2\log \frac{1}{\pi(B_{L_1(S)}(f,\varepsilon_1))}}{n}} + 2(\varepsilon_1 - 1/n) \right\} &\leq \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2\log \frac{1}{\pi(B_{L_1(S')}(f,\varepsilon))}}{n}} + 2\varepsilon \right\} \\ &\leq \inf_{\varepsilon_2 > 0} \left\{ \sqrt{\frac{2\log \frac{1}{\pi(B_{L_1(S')}(f,\varepsilon_2))}}{n}} + 2(\varepsilon + 1/n) \right\}. \end{split}$$

This sandwich bound implies that

$$\left|\inf_{\varepsilon>1/n}\left\{\sqrt{\frac{2\log\frac{1}{\pi(B_{L_1(S')}(f,\varepsilon))}}{n}}+2\varepsilon\right\}-\inf_{\varepsilon>1/n}\left\{\sqrt{\frac{2\log\frac{1}{\pi(B_{L_1(S)}(f,\varepsilon))}}{n}}+2\varepsilon\right\}\right|\leq \frac{2}{n}.$$

Combining the above bound with (23) we prove the bounded difference property (22). Thus, by McDiarmid's inequality (Lemma 10), from the bounded difference property (22), with probability at least $1 - \delta$,

$$\begin{split} \sup_{f \in \mathcal{F}} \left\{ (\mathbb{P} - \mathbb{P}_n) \ell(f; z) - \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2 \log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f, \varepsilon))}}{n}} + 2\varepsilon \right\} - \sqrt{\frac{4}{n}} \right\} \\ \leq & \mathbb{E}_z \sup_{f \in \mathcal{F}} \left\{ (\mathbb{P} - \mathbb{P}_n) \ell(f; z)) - \inf_{\varepsilon > 1/n} \left\{ \sqrt{\frac{2 \log \frac{1}{\pi(B_{L_1(\mathbb{P}_n)}(f, \varepsilon))}}{n}} + 2\varepsilon \right\} - \sqrt{\frac{4}{n}} \right\} + 3\sqrt{\frac{\log \frac{1}{\delta}}{2n}} \\ \leq & 3\sqrt{\frac{\log \frac{1}{\delta}}{2n}}. \end{split}$$

Finally, we convert the stronger $L_1(\mathbb{P}_n)$ bound to the weaker $L_2(\mathbb{P}_n)$ bound we stated in the main paper: combining the above $L_1(\mathbb{P}_n)$ bound with the fact $L_1(\mathbb{P}_n) \leq L_2(\mathbb{P}_n)$, we have that with probability at least $1 - \delta$, uniformly over every $h \in \mathcal{H}$,

$$(\mathbb{P} - \mathbb{P}_n)\ell(f; z) \le \inf_{\varepsilon > \sqrt{1/n}} \left\{ 2\varepsilon + \sqrt{\frac{2\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\varepsilon))}}{n}} + \frac{4}{n} + 3\sqrt{\frac{\log \frac{1}{\delta}}{2n}} \right\}.$$

B.2 Proof for Theorem 2 (Generic Chaining Upper and Lower Bounds)

Theorem 2 is a corollary of established matching upper and lower bounds for Gaussian Processes Fernique (1975); Talagrand (1987). Note that generic chaining have several equivament formulations, and the one closest to our purpose is through majorizing measure. On the upper bound side, we extend it to pointwise generalization bound; on the lower bound side, we re-derive it to empirical processes.

B.2.1 BACKGROUND ON GAUSSIAN PROCESSES

We now recall several key results from a series of seminal papers by Talagrand, Fernique, and others, which introduces the majorizing-measure formulation of the generic chaining framework.

The statements below follow Talagrand's original exposition [Talagrand, 1987, "Regularition of Gaussian Processes," Acta Mathematica, 1987] closely, with only minor adjustments in notation and wording.

A centered Gaussian random variable X is a real-valued measurable function on the outcome space such that the law of X has density

$$(2\pi\sigma^2)^{-1/2} \exp\left(-\frac{x^2}{2\sigma^2}\right).$$

The law of X is thus determined by $\sigma = (\mathbb{E}[X^2])^{1/2}$. If $\sigma = 1, X$ is called *standard normal*.

A Gaussian process is a family $\{X_t\}_{t\in T}$ of random variables indexed by some set T, such that every finite linear combination $\sum_{j=1}^k \alpha_j X_{t_j}$ is Gaussian. Its covariance function

$$\Gamma(u,v) = \mathbb{E}[X_u X_v], \quad (u,v) \in T \times T,$$

determines $\mathbb{E}\left[\sum_{j} \alpha_{j} X_{t_{j}}\right]^{2}$, and hence the joint law of $\{X_{t}\}_{t \in T}$. On the index set T, consider the pseudo-distance ϱ given by

$$\varrho(u,v) = \sqrt{\mathbb{E}[(X_u - X_v)^2]}.$$

Gaussian processes are thus a very rigid class of stochastic processes, with exceptionally nice properties that have been fully developed in the literature.

Fernique Fernique (1975) proved in 1975 the following result.

Lemma 6 (Upper Bound of Gaussian Processes via Majorizing Measure, Fernique (1975)) For any Gaussian process $(X_t)_{t\in T}$, and any prior π on T, we have

$$\mathbb{E}\left[\sup_{t\in T}X_t\right] \leq C\sup_{x\in T}\int_0^\infty \sqrt{\log\frac{1}{\pi(B_\varrho(x,\epsilon))}}d_\epsilon,$$

where C > 0 is an absolute constant.

A prior π that makes the right hand side in Lemma 6 finite is called a *majorizing measure*. Fernique conjectured as early as 1974 that the existence of majorizing measures might characterize the boundedness of Gaussian processes. He proved a number of important partial results, and his determination eventually motivated the Talagrand to attack the problem in 1987. In Talagrand (1987) Talagrand proved that the integral in Lemma 6 is tight up to absolute constants. We define the following minor measurability condition:

Assumption 1 (Minor Measurability Condition) The Gaussian process satisfies $\sup_{t \in D} |X_t| < \infty$ almost surely for each countable subset $D \subseteq T$.

Lemma 7 (Lower Bound of Gaussian Processes via Majorzing Measure, Talagrand (1987))

For every Gaussian process that satisfies a mild measurability condition (Assumption 1), there exists a prior π on T such that

$$\mathbb{E}\left[\sup_{x\in T} X_t\right] \ge c \sup_{x\in T} \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho}(x,\epsilon))}} d\epsilon,$$

where c > 0 is an absolute constant.

Thus the pointwise dimension integral gives a complete characterization to the supremum of Gaussian process. We now use the upper and lower bounds in Lemma 6 and Lemma 7 to prove Theorem 2.

B.2.2 PROOF OF THEOREM 2

 I: Proof of Upper Bound. We begin by citing a general principle for converting (localized) uniform convergence guarantees into pointwise generalization bounds. This conversion, introduced as "uniform localized convergence" principle in Xu & Zeevi (2025), provides a direct mechanism for obtaining the type of pointwise generalization bounds central to our work. We state this result as "uniform pointwise convergence" principle.

Lemma 8 ("Uniform Pointwise Convergence" Principle Xu & Zeevi (2025):) (A Generic Conversion from Uniform Convergence to Pointwise Generalization.) For a function class \mathcal{F} and functional $d: \mathcal{F} \to (0, R]$, assume there is a function $\psi(r; \delta)$, which is non-decreasing with respect to r, non-increasing with respect to δ , and satisfies that $\forall \delta \in (0, 1), \forall r \in [0, R]$, with probability at least $1 - \delta$,

$$\sup_{f \in \mathcal{F}: d(f) \le r} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) \le \psi(r; \delta). \tag{24}$$

Then, given any $\delta \in (0,1)$ and $r_0 \in (0,R]$, with probability at least $1-\delta$, uniformly over all $f \in \mathcal{F}$,

$$(\mathbb{P} - \mathbb{P}_n)\ell(f; z) \le \psi\left(\max\{2d(f), r_0\}; \delta\left(\log_2 \frac{2R}{r_0}\right)^{-1}\right). \tag{25}$$

The proof of the upper bound in Theorem 2 consists of three steps: 1. Bounding Empirical Process by Gaussian Process; 2. Applying Integral Upper Bound; 3. Generic Conversion to Pointwise Generalization Bound.

Step 1: Bounding Empirical Process by Gaussian Process. We start with the classical result that Gaussian complexity of a function class is an upper bound for expected uniform convergence, see Definition 2 and Lemma 11 in the auxiliary lemma part for this classical result. To be specific, by Lemma 11 (and $\sqrt{2\pi} < 3$) we have

$$\mathbb{E}_{z} \left[\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_{n}) \ell(f; z) \right] \leq 3 \mathbb{E}_{g, z} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} g_{i} \ell(f; z_{i}) \right],$$

where g_i are i.i.d. standard Gaussian variables. Applying Mcdiarmid's inequality (Lemma 10) twice, we have that with probability at least $1 - \delta$, the following two inequality simultaneously hold:

$$\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) - \mathbb{E}_z \left[\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) \right] \le \sqrt{\frac{\log \frac{2}{\delta}}{2n}},$$

$$\mathbb{E}_{g, z} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i \ell(f; z_i) \ell(f; z) \right] - \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i \ell(f; z_i) \ell(f; z) \le \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

Combining the above three inequalities we have

$$\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) \le 3 \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i \ell(f; z_i) + \sqrt{\frac{8 \log \frac{2}{\delta}}{n}},$$

taking expectation for $\{g_i\}_{i=1}^n$ on both sides, we obtain

$$\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) \le 3\mathbb{E}_g \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i \ell(f; z_i) \right] + \sqrt{\frac{8 \log \frac{2}{\delta}}{n}}, \tag{26}$$

Step 2: Applying Integral Upper Bound. Applying Lemma 6 to the Gaussian process $\frac{1}{n}\sum_{i=1}^n g_i\ell(f;z_i)$ (with fixed $\{z_i\}$), we have that for all data-dependent prior π ,

$$\mathbb{E}_g \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i \ell(f; z_i) \right] \le \frac{C_1}{\sqrt{n}} \sup_{f \in \mathcal{F}} \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon, \tag{27}$$

where $C_1 > 0$ is an absolute constant.

 Step 3: Generic Conversion to Pointwise Generalization Bound. Combining (26) and (27), we have: for all $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) \leq \frac{3C_1}{\sqrt{n}} \sup_{f \in \mathcal{F}} \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon + \sqrt{\frac{8\log \frac{2}{\delta}}{n}},$$

which implies that $\forall \delta \in (0,1)$ and $\forall r \in (0,n]$, with probability at least $1-\delta$

$$\sup_{f:d(f) < r} (\mathbb{P} - \mathbb{P}_n)\ell(f; z) \le C_2 \sqrt{\frac{r}{n}} + \sqrt{\frac{8\log\frac{2}{\delta}}{n}},\tag{28}$$

where the functional $d: \mathcal{F} \to [0, n]$ is defined by

$$d(f) = \min \left\{ \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\epsilon))}} d\epsilon, \sqrt{n} \right\}^2,$$

and $C_2 = \max\{3C_1, 1\}$ is an absolute constant $(C_1 \text{ controls } (28) \text{ for } r < n \text{ and } 1 \text{ controls } (28) \text{ for } r = n)$. The inequality (28) is precisely the condition (24) in the generic conversion provided in Lemma 8. Thus applying Lemma 8 we have the pointwise generalization bound: for any $\delta \in (0,1)$, and $r_0 = 1/n$, with probability at least $1 - \delta$, uniformly over all $f \in \mathcal{F}$,

$$\begin{split} (\mathbb{P} - \mathbb{P}_n) \ell(f; z) \leq & \frac{C_2}{\sqrt{n}} \max \left\{ 2 \min \left\{ \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon, \sqrt{n} \right\}, \frac{1}{n} \right\} + \sqrt{\frac{8 \log \frac{\log_2(4n^2)}{\delta}}{n}} \\ \leq & \frac{2C_2}{\sqrt{n}} \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon + \frac{C_2}{n^{1.5}} + \sqrt{\frac{8 \log \frac{\log_2(4n^2)}{\delta}}{n}} \\ \leq & C \left(\frac{1}{\sqrt{n}} \int_0^\infty \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon + \sqrt{\frac{\log \frac{\log(2n)}{\delta}}{n}} \right), \end{split}$$

where C > 0 is an absolute constant.

II: Proof of Lower Bound. We use the classical result that the expected uniform convergence is lower bounded by Gaussian complexity of the centered class, up to a $\sqrt{\log n}$ factor, see Definition 2 and Lemma 12 in the auxiliary lemma part for this classical result. To be specific, by Lemma 12 we have that

$$\mathbb{E}_{z} \left[\sup_{f \in \mathcal{F}} (\mathbb{P} - \mathbb{P}_{n}) \ell(f; z) \right] \geq \frac{c_{1}}{\sqrt{\log n}} \mathbb{E}_{g, z} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} g_{i} (\ell(f; z_{i}) - \mathbb{E}_{z}[\ell(f; z)]) \right] \\
\geq \frac{c_{1}}{\sqrt{\log n}} \mathbb{E}_{g, z} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} g_{i} \ell(f; z_{i}) - \left| \frac{1}{n} \sum_{i=1}^{n} g_{i} \right| \cdot \sup_{\mathcal{F}} \mathbb{E}[\ell(f; z)] \right] \\
= \frac{c_{1}}{\sqrt{\log n}} \mathbb{E}_{g, z} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} g_{i} \ell(f; z_{i}) \right] - \frac{c_{1}}{\sqrt{\log n}} \sqrt{\frac{2}{\pi n}} \sup_{\mathcal{F}} \mathbb{E}[\ell(f; z)], \quad (29)$$

where $c_1 > 0$ is an absolute constant, and the equality use the fact that $\mathbb{E}[Y] = \sqrt{\frac{2}{\pi n}}$ for $Y \sim N(0, 1/n)$.

Now applying Lemma 7 to lower bounding the Gaussian process $\frac{1}{n} \sum_{i=1}^{n} g_i \ell(f; z_i)$ by the integral, we have for any $\{z_i\}_{i=1}^n$,

$$\mathbb{E}_g \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i \ell(f; z_i) \right] \ge c_2 \inf_{\pi} \sup_{f \in \mathcal{F}} \int_0^{\infty} \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon,$$

taking expectation on both side yields

$$\mathbb{E}_{g,z} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} g_i \ell(f; z_i) \right] \ge c_2 \mathbb{E} \inf_{\pi} \sup_{f \in \mathcal{F}} \int_0^{\infty} \sqrt{\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(f, \epsilon))}} d\epsilon.$$
 (30)

Combining (29) and (30), we have that there exist absolute constants c, c' > 0 such that

$$\mathbb{E}\left[\sup_{f\in\mathcal{F}}(\mathbb{P}-\mathbb{P}_n)\ell(f;z)\right] \geq \frac{c}{\sqrt{n\log n}}\mathbb{E}\inf_{\pi}\sup_{f\in\mathcal{F}}\int_0^{\infty}\sqrt{\log\frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\epsilon))}}d\epsilon - \frac{c'\sup_{\mathcal{F}}\mathbb{E}[\ell(f;z)]}{\sqrt{n\log n}}.$$

This inequality implies the following result

$$\mathbb{E}\left[\inf_{\pi}\sup_{f\in\mathcal{F}}\left((\mathbb{P}-\mathbb{P}_n)\ell(f;z)-\frac{c}{\sqrt{n\log n}}\int_0^{\infty}\sqrt{\log\frac{1}{\pi(B_{\varrho_{n,\ell}}(f,\varepsilon))}}d\varepsilon\right)+\frac{c'\sup_{\mathcal{F}}\mathbb{E}[\ell(f;z)]}{\sqrt{n\log n}}\right]\geq 0,$$

B.3 AUXILIARY LEMMAS FOR SECTION 2

Lemma 9 (PAC-Bayes Bound (Catoni, 2003); see also Theorem 2.1 in Alquier et al. (2024))

Let π be any prior on a hypothesis class \mathcal{F} , and let $\ell \colon \mathcal{F} \times \mathcal{Z} \to [0, B]$ be a bounded loss. Fix confidence $\delta \in (0, 1)$ and sample size n. Then for every $\eta > 0$, with probability at least $1 - \delta$ over n i.i.d. draws $z_1, \ldots, z_n \sim P$, for every distribution μ on \mathcal{F} simultaneously,

$$(\mathbb{P} - \mathbb{P}_n)\ell(f; z) \leq \sqrt{\frac{B^2(\mathrm{KL}(\mu, \pi) + \log \frac{1}{\delta})}{8n}}.$$

Lemma 10 (McDiarmid's inequality (McDiarmid, 1998)) Suppose that $z_1, ..., z_n \in \mathcal{X}$ are independent, and $h: \mathcal{Z}^n \to \mathbb{R}$. Let $c_1, ..., c_n$ satisfy

$$\sup_{z_1,...,z_n,z_i'} |h(z_1,...,z_{i-1},z_i,z_{i+1},...,z_n) - h(z_1,...,z_{i-1},z_i',z_{i+1},...,z_n)| \le c_i,$$

for i = 1, ..., n. Then

$$\Pr(h(z_1,\dots,z_n) - \mathbb{E}[h(z_1,\dots,z_n)] \ge t) \le \exp\left(\frac{-2t^2}{\sum_{i=1}^n c_i^2}\right).$$

Definition 2 (Rademacher and Gaussian complexities) For a function class \mathcal{F} that consists of mappings from \mathcal{Z} to \mathbb{R} , define the Rademacher complexity of \mathcal{F} as

$$R_n(\mathcal{F}) := \mathbb{E}\left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \xi_i f(z_i)\right],$$

where ξ_i are i.i.d. Rademacher variables; and define the Gaussian complexity of \mathcal{F} as

$$G_n(\mathcal{F}) := \mathbb{E}\left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n g_i f(z_i)\right],$$

where g_i are i.i.d. standard Gaussian variables.

Lemma 11 (Upper Bounds with Rademacher and Gaussian Complexities, Lemma 7.4 in Van Handel (2014)) For any function class \mathcal{F} that consists of mappings from \mathcal{Z} to \mathbb{R} , we have

$$\mathbb{E}\left[\sup_{f\in\mathcal{F}}(\mathbb{P}-\mathbb{P}_n)f(z)\right] \leq 2R_n(\mathcal{F}) \leq \sqrt{2\pi}G_n(\mathcal{F}).$$

Lemma 12 (Lower Bounds with Rademacher and Gaussian Complexities) For any function class \mathcal{F} that consists of mappings from \mathcal{Z} to \mathbb{R} , defined its centered class $\tilde{\mathcal{F}}$ as $\{f - \mathbb{E}[f(z)] : f \in \mathcal{F}\}$. We have

$$\mathbb{E}\left[\sup_{f\in\mathcal{F}}(\mathbb{P}-\mathbb{P}_n)f(z)\right] \ge \frac{1}{2}R_n(\tilde{\mathcal{F}}) \ge \frac{c}{\sqrt{\log n}}G_n(\tilde{\mathcal{F}}),$$

where c > 0 is an absolute constant.

Proof of Lemma 12: Both the fact that uniform convergence admit a lower bound in terms of the Rademacher complexity of the centered class, and the result that Rademacher complexity itself is bounded below by Gaussian complexity up to a factor of $\sqrt{\log n}$, are classical and admit simple proofs. For a full proof of the first inequality, see Theorem 14.3 in Rinaldo & Yan (2016); for a reference and proof sketch of the second inequality, see Problem 7.1 in Van Handel (2014).

C PROOFS FOR DEEP NEURAL NETWORKS AND RIEMANNIAN DIMENSION (SECTION 3)

C.1 Proof of Lemma 1 in Section 3.1

We start with the telescoping decomposition presented in the main paper, which serves as a non-perturvative replacement of conventional Taylor expansion, where in each summand the only difference lie in W_l' and W_l .

$$\begin{split} &F_L(W',X) - F_L(W,X) \\ &= \sum_{l=1}^L [\underbrace{\sigma_L(W'_L \cdots W'_{l+1}}_{\text{controled by} M_{l \to L}} \underbrace{\sigma_l}_{\text{by1}} (W'_l \underbrace{F_{l-1}(W,X)}_{\text{learned feature}})) - \sigma_L(W'_L \cdots W'_{l+1} \sigma_l (W_l \underbrace{F_{l-1}(W,X)}_{\text{learned feature}}))], \end{split}$$

Applying Cauchy-Schwartz inequality to the above identity, we have

$$||F(W',X) - F(W,X)||_{\mathbf{F}}^2$$
 (31)

$$\leq \sum_{l=1}^{L} L \|\sigma_{L}(W'_{L} \cdots W'_{l+1} \sigma_{l}(W'_{l} F_{l-1}(W, X))) - \sigma_{L}(W'_{L} \cdots W'_{l+1} \sigma_{l}(W_{l} F_{l-1}(W, X)))\|_{\mathbf{F}}^{2}$$
(32)

By the definition of local Lipchitz constant in Section 3, for all $W' \in B_{\varrho_n}(W, \varepsilon)$,

$$\|\sigma_{L}(W'_{L}\cdots W'_{l+1}\sigma_{l}(W'_{l}F_{l-1}(W,X))) - \sigma_{L}(W'_{L}\cdots W'_{l+1}\sigma_{l}(W_{l}F_{l-1}(W,X)))\|_{\mathbf{F}}$$

$$\leq M_{l\to L}[W,\epsilon]\|\sigma_{l}(W'_{l}F_{l-1}(W,X)) - \sigma_{l}(W_{l}F_{l-1}(W,X))\|_{\mathbf{F}}.$$
(33)

Because the activation function σ_l is 1—Lipchitz for each column, we have

$$\|\sigma_l(W_l'F_{l-1}(W,X)) - \sigma_l(W_lF_{l-1}(W,X))\|_{\mathbf{F}} \le \|(W_l' - W_l)F_{l-1}(W,X)\|_{\mathbf{F}}.$$
 (34)

 Combining (31) (33) and (34), we prove that

$$\|F(W',X) - F(W,X)\|_{\mathbf{F}}^2 \le \sum_{l=1}^L L \cdot M_{l \to L}[W,\epsilon]^2 \cdot \|(W'_l - W_l)F_{l-1}(W,X)\|_{\mathbf{F}}^2.$$

C.2 METRIC DOMINATION LEMMA

Our non-perturbative expansion facilitates bounding the pointwise dimension of complex geometries via metric comparison. By constructing a simpler, dominating metric (i.e., one that is pointwise larger), we establish that the pointwise dimension of the original geometry is upper bounded by that of this new, more structured geometry. This "enlargement" for analytical tractability, a concept with roots in comparison geometry and majorization principles, is operationalized in Lemma 13.

Lemma 13 (Metric Domination Lemma) For two metrics ϱ_1, ϱ_2 defined on \mathbb{R}^p , if $\varrho(W', W) \leq \varrho'(W', W)$ for all $W' \in B_{\varrho_2}(W, \epsilon)$, then for any prior $\pi \in \Delta(\mathbb{R}^p)$ and any $\epsilon > 0$, we have

$$\log \frac{1}{\pi(B_{\varrho_1}(W,\varepsilon))} \le \log \frac{1}{\pi(B_{\varrho_2}(W,\epsilon))}.$$

Proof of Lemma 13: Because $\varrho_1(W',W) \leq \varrho_2(W',W)$ for all $W' \in B_{\varrho_2}(W,\epsilon)$, we have that $B_{\varrho_1}(W,\epsilon) \supseteq B_{\varrho_2}(W,\epsilon)$.

So for any prior π on \mathbb{R}^p , monotonicity of measures gives

$$\pi(B_{\rho_1}(W,\epsilon)) \geq \pi(B_{\rho_2}(W,\epsilon)),$$

this implies

$$\log \frac{1}{\pi(B_{\varrho_1}(W,\varepsilon))} \leq \log \frac{1}{\pi(B_{\varrho_2}(W,\epsilon))}.$$

We then state an extension of the metric domination lemma, which turns pointwise dimension in a high-dimensional space into a lower-dimensional sub-space.

Lemma 14 (Sub-Space Metric Domination Lemma) Given a metrics ϱ_1 defined on \mathbb{R}^p a subspace $\mathcal{V} \subseteq \mathbb{R}^p$, and a metrics ϱ_2 defined on \mathcal{V} . Define the orthogonal projector to sub-space \mathcal{V} as $\mathcal{P}_{\mathcal{V}}(W) := \arg\min_{\tilde{W} \in \mathcal{V}} \|\tilde{W} - W\|_2$. If there exists $\epsilon_1 \in (0, \epsilon)$ such that for every $W' \in \mathcal{V}$,

$$(\varrho_1(W', W))^2 \le (\varrho_2(W', \mathcal{P}_{\mathcal{V}}(W)))^2 + \epsilon_1^2,$$
 (35)

then for any prior $\pi \in \Delta(\mathcal{V})$, we have

$$\log \frac{1}{\pi(B_{\varrho_1}(W,\varepsilon))} \le \log \frac{1}{\pi(B_{\varrho_2}(\mathcal{P}_{\mathcal{V}}(W), \sqrt{\epsilon^2 - \epsilon_1^2}))}.$$
 (36)

Proof of Lemma 14: By the condition (35), we know

$$B_{\varrho_1}(W,\varepsilon) \supseteq B_{\varrho_1}(W,\varepsilon) \cap \mathcal{V} \supseteq B_{\varrho_2}(\mathcal{P}_{\mathcal{V}}(W), \sqrt{\epsilon^2 - \epsilon_1^2}),$$

and this gives the desired conclusion (36) in Lemma 14.

We also show that dominating matrix always has larger eigenvalues, which will be useful in our proof.

Lemma 15 (Eigenvalue monotonicity under Loewner order) Let $A, B \in \mathbb{R}^{p \times p}$ be positive-semidefinite (PSD) matrices. If $A \succeq B$ (i.e. A - B is PSD), then for every $k \in \{1, \ldots, p\}$

$$\lambda_k(A) \geq \lambda_k(B),$$

where $\lambda_1 \geq \cdots \geq \lambda_p$ denotes the eigenvalues in nonincreasing order.

Proof of Lemma 15: The Courant–Fischer–Weyl max-min characterization Wikipedia contributors (b) states that for any Hermitian (i.e., symmetric for real matrices studying here) matrix M,

$$\lambda_k(M) = \max_{\substack{\mathcal{V} \subseteq \mathbb{R}^p \\ \dim \mathcal{V} = k}} \min_{\substack{W \in \mathcal{V} \\ W \neq 0}} \frac{W^\top M W}{\|W\|^2}.$$

Because $A \succeq B$ implies $W^{\top}BW \leq W^{\top}AW$ for every vector W, we have, for every k-dimensional subspace \mathcal{V} ,

$$\min_{W \in \mathcal{V} \backslash \{0\}} \frac{W^\top B W}{\|W\|^2} \ \leq \ \min_{W \in \mathcal{V} \backslash \{0\}} \frac{W^\top A W}{\|W\|^2}.$$

Taking the maximum over all such subspaces yields

$$\lambda_k(B) \leq \lambda_k(A),$$

as desired.

C.2.1 AUXILIARY LEMMAS FOR VOLUMETRIC ARGUMENTS

The following result on volumes of d-dimensional ℓ_2 ball with radius r is classical (e.g., see the Wikipedia page Wikipedia contributors (d)).

Lemma 16 (Volumes of ℓ_2 **Ball, Wikipedia contributors (d))** *Let* d *be a positive integer and* r > 0. *Define*

$$B_2(r) = \{ x \in \mathbb{R}^d : ||x||_2 \le r \}.$$

Then its volume is

$$\operatorname{Vol}(B_2(r)) = \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)} r^d.$$

We now state the result on the volume of an ellipsoid, which is also classical. We give a proof here.

Lemma 17 (Volume of a Σ **-Ellipsoid)** *Let* $\Sigma \in \mathbb{R}^{d \times d}$ *be symmetric positive definite with eigenvalues* $\lambda_1, \ldots, \lambda_d > 0$. *For any* r > 0, *define the ellipsoid*

$$E := \{ w \in \mathbb{R}^d : w^{\top} \Sigma w \le r^2 \}.$$

Then its d-dimensional volume is

$$\operatorname{Vol}(E) \ = \ \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)} \ r^d \ \Big(\prod_{i=1}^d \lambda_i\Big)^{-\frac{1}{2}}.$$

Proof of Lemma 17: Denote $\Sigma^{1/2}$ to be the square root of Σ : the unique positive semidefinite matrix satisfying $\Sigma^{1/2}\Sigma^{1/2} = \Sigma$ (see, e.g., Wikipedia contributors (c)). Define the linear map

$$T: \mathbb{R}^d \to \mathbb{R}^d, \qquad z = T(w) := \Sigma^{1/2} w.$$

For $w \in E$ we have $w^{T} \Sigma w \leq r^{2}$, hence $||z||_{2}^{2} = w^{T} \Sigma w \leq r^{2}$. Thus

$$T(E) = B_2(r) := \{ z \in \mathbb{R}^d : ||z||_2 \le r \},$$

the Euclidean ball of radius r. By the change-of-variables theorem in multivariate calculus Wikipedia contributors (a), the linear map T scales Lebesgue measure by $|\det T|$. Thus

$$Vol(E) = |\det T|^{-1} Vol(T(E)) = (\det \Sigma)^{-1/2} Vol(B_2(r)),$$

where the second identity holds since $T = \Sigma^{1/2}$,

$$|\det T| = \det(\Sigma^{1/2}) = (\det \Sigma)^{1/2}.$$

The standard formula, Lemma 16, gives

$$\operatorname{Vol}(B_2(r)) = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)} r^d.$$

This yields the claimed expression

$$\operatorname{Vol}(E) \ = \ (\det \Sigma)^{-1/2} \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)} \, r^d \ = \ \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)} \, r^d \, \left(\prod_{i=1}^d \lambda_i \right)^{-\frac{1}{2}}.$$

C.3 POINTWISE DIMENSION BOUND WITH APPROXIMATE EFFECTIVE SUB-SPACE

Set Up of Approximate Effective Sub-Space Let the $p \times p$ PSD matrix Σ have the eigendecomposition $\Sigma = U \operatorname{diag}(\lambda_1, \cdots, \lambda_p) U^{\top}$, where $\lambda_1, \cdots, \lambda_p$ are in non-increasing order, and $U = [u_1; \cdots; u_p]$. Use r_{eff} as the abbreviation of $r_{\mathrm{eff}}(\Sigma, R, \epsilon)$, and $\mathcal V$ as an abbreviation of $\mathcal V_{\mathrm{eff}}(\Sigma, R, \epsilon)$ —the true effective sub-space (i.e., top- r_{eff} eigenspace).

Assume there is another r-dimensional sub-space $\bar{\mathcal{V}}$. We will show that if $\bar{\mathcal{V}}$ approximates \mathcal{V} , then using a prior supported on $\bar{\mathcal{V}}$ still yields a valid effective-dimension bound. This observation underpins the hierarchical covering argument in Theorem 3. For a self-contained introduction to subspaces (collectively known as the Grassmannian) and their frame parameterizations (the Stiefel manifold), see Section D.1, where we translate algebraic and differential-geometric insights into machine learning terminology.

Motivation of Approximate Effective Sub-Space. We can view the orthogonal projector to a subspace as a matrix (see the definition via the Stiefel parameterization in (58)), which is consistent with the earlier operator notation characterized by ℓ_2 -distance in Lemma 14. Now define the projected metric $\varrho_{\Sigma}^{\bar{\nu}}$ as

$$\varrho_{\Sigma}^{\vec{\mathcal{V}}}(W_1, W_2) = \sqrt{(\mathcal{P}_{\vec{\mathcal{V}}}(W_1) - \mathcal{P}_{\vec{\mathcal{V}}}(W_1))^{\top} \Sigma (\mathcal{P}_{\vec{\mathcal{V}}}(W_1) - \mathcal{P}_{\vec{\mathcal{V}}}(W_2))} = \sqrt{(W_1 - W_2) \mathcal{P}_{\vec{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\vec{\mathcal{V}}}(W_1 - W_2)}.$$

By the sub-space metric dominance lemma (Lemma 14), if $\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}}$ approximates Σ , we can use prior over $\bar{\mathcal{V}}$ to bound the pointwise dimension and achieve dimension reduction.

We will require the following approximation error condition:

$$\|\Sigma^{\frac{1}{2}}(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{op}} \le \frac{\sqrt{n}\epsilon}{4R}.$$
(37)

In Section D, we systematically study the ellipsoidal covering of Grassmannian, and establish that we can *always* find $\bar{\mathcal{V}}$ that approximates \mathcal{V} to the desired precision, with an additional covering cost of the Grassmannain bound in the Riemannain Dimension. This generalizes the canonical projection metric between sub-spaces into ellipsoidal set-up.

Effective Dimension Bound for Approximate Effective Sub-Space. We now present the lemma that establish effective dimension bound using prior supported on approximate effective sub-space $\bar{\mathcal{V}}$ (not necessarily the true effective sub-space $\mathcal{V}_{\text{eff}}(\Sigma, R, \epsilon)$). We state the main result of this subsection (Lemma 2) in the main paper).

Consider a manifold (W, Σ) $W \subseteq B_2(R) \subset \mathbb{R}^p$ and G(W) defines the metric tensor, and a given r-dimensional sub-space $\bar{\mathcal{V}} \subseteq \mathbb{R}^p$. The prior $\pi_{\bar{\mathcal{V}}} = \mathrm{Unif}(B_2(1.58R) \cap \bar{\mathcal{V}})$ satisfies that, uniformly over all (W, ϵ) such that $\bar{\mathcal{V}}$ approximate $\mathcal{V}_{\mathrm{eff}}(G(W), R, \epsilon)$ to the precision (37),

$$\log \frac{1}{\pi_{\bar{\mathcal{V}}}(B_{\varrho_G(W)}(W,\sqrt{n}\varepsilon))} \, \leq \, \frac{1}{2} \sum_{k=1}^{r_{\mathrm{eff}}(G(W),R,\varepsilon)} \log \left(\frac{40R^2\lambda_k}{n(\varepsilon^2-\epsilon_1^2)} \right) = d_{\mathrm{eff}}(\Sigma,\sqrt{5}R,\epsilon).$$

Proof of Lemma 2: Given a PSD matrix Σ with eigenvalues $\lambda_1 \geq \cdots \lambda_d$, denote $r_{\text{eff}} = r_{\text{eff}}(\Sigma, R, \epsilon)$, and the projected metric $\varrho_{\Sigma}^{\bar{V}}$ on \bar{V} :

$$\varrho_{\Sigma}^{\vec{\mathcal{V}}}(W_1, W_2) = \sqrt{(W_1 - W_2)^{\top} \mathcal{P}_{\vec{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\vec{\mathcal{V}}}(W_1 - W_2)}.$$

Since V is the top- r_{eff} eigenspace of Σ , by the elementary property of eigendecomposition we have that

$$\Sigma = \mathcal{P}_{\mathcal{V}}^{\top} \Sigma \mathcal{P}_{\mathcal{V}} + \mathcal{P}_{\mathcal{V}_{\perp}}^{\top} \Sigma \mathcal{P}_{\mathcal{V}_{\perp}}$$

$$\leq \mathcal{P}_{\mathcal{V}}^{\top} \Sigma \mathcal{P}_{\mathcal{V}} + \lambda_{r_{\text{eff}}+1} \cdot \mathcal{P}_{\mathcal{V}_{\perp}}^{\top} \mathcal{P}_{\mathcal{V}_{\perp}},$$
(38)

where \mathcal{V}_{\perp} is orthogonal complement of \mathcal{V} . It is also straightforward to see

$$\mathcal{P}_{\mathcal{V}}^{\top} \Sigma \mathcal{P}_{\mathcal{V}} \leq 2 \mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}} + 2 (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})^{\top} \Sigma (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}}). \tag{39}$$

Combining (38) and (39), we have the fundamental loewner order inequality

$$\Sigma \leq 2\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}} + 2(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})^{\top} \Sigma (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}}) + \lambda_{r_{\text{eff}} + 1} \cdot \mathcal{P}_{\mathcal{V}_{\perp}}^{\top} \mathcal{P}_{\mathcal{V}_{\perp}}. \tag{40}$$

In order to apply the sub-space metric domination lemma (Lemma 14), we hope to bound $||W' - W||_2^2$ and apply that bound to the two last reminder terms in the right hand side of (40).

To bound $||W' - W||_2^2$, we firstly state the following lemma on the eigenvalue of $\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}}$, whose proof is deferred until after the current proof.

Lemma 18 (Eigenvalue Bound for Projected Metric Tensor) Assume V is the top-r eigenspace of Σ , then for a r-dimensional sub-space \bar{V} we have that for $k = 1, 2, \dots, r$,

$$\lambda_k \ge \lambda_k (\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}}) \ge \lambda_k / 2 - \|\Sigma^{\frac{1}{2}} (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{op}}^2.$$

For every $W'\in B_{\varrho_{\mathcal{X}}^{\mathcal{Y}}}(\mathcal{P}_{\mathcal{V}}(W),\sqrt{n}\varepsilon/4)$, we have $\forall k=1,\cdots,r_{\mathrm{eff}}$,

$$\|W' - \mathcal{P}_{\bar{\mathcal{V}}}(W)\|_{2}^{2} \leq \frac{(W' - \mathcal{P}_{\bar{\mathcal{V}}}(W))\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}}(W' - \mathcal{P}_{\bar{\mathcal{V}}}(W))}{\lambda_{r_{\text{eff}}}(\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}})} \leq \frac{n\epsilon^{2}}{16\lambda_{r_{\text{eff}}}(\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}})}$$

$$\leq \frac{n\epsilon^{2}}{8\lambda_{r_{\text{eff}}} - 16\|\Sigma^{\frac{1}{2}}(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{on}}^{2}} \leq \frac{1}{3}R^{2}, \tag{41}$$

where the last inequality uses $\lambda_{r_{\rm eff}} \geq \frac{n\epsilon^2}{2R^2}$ and the approximation error condition (37). On the other hand, we have that $\|W\|_2^2 \leq R^2$, so that for every $W' \in B_{\varrho_{\Sigma}^{\mathcal{V}}}(\mathcal{P}_{\mathcal{V}}(W), \sqrt{n}\varepsilon/2)$

$$||W' - W||_2^2 = ||W' - \mathcal{P}_{\bar{\mathcal{V}}}(W)||_2^2 + ||\mathcal{P}_{\bar{\mathcal{V}}_{\perp}}(W)||_2^2 \le \frac{4}{3}R^2.$$

From the fundamental loewner order inequality (40), we establish the desired metric domination condition: for all $W' \in B_{\rho_{\chi}^{V}}(W, \sqrt{n}\epsilon/3)$ and $W \in B_{2}(R)$,

$$(W' - W)^{\top} \Sigma (W' - W)$$

$$\leq (W' - W)^{\top} (2\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}}) (W' - W) + (2\|\Sigma^{\frac{1}{2}} (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{op}}^{2} + \lambda_{r_{\text{eff}}+1}) \|W' - W\|_{2}^{2}$$

$$\leq 2\varrho_{\Sigma}^{\bar{\mathcal{V}}} (W', \mathcal{P}_{\bar{\mathcal{V}}}(W))^{2} + \frac{5n\epsilon^{2}}{6}.$$

Now we can apply the sub-space metric domination lemma (Lemma 14) and obtain: for any $\pi \in \Delta(\bar{\mathcal{V}})$,

$$\log \frac{1}{\pi \left(B_{\varrho_{\Sigma}}(W, \sqrt{n}\varepsilon)\right)} \le \log \frac{1}{\pi \left(B_{\sqrt{2}\varrho_{\Sigma}^{\widetilde{V}}}(W, \sqrt{n}\epsilon/\sqrt{6})\right)} \le \log \frac{1}{\pi \left(B_{\varrho_{\Sigma}^{\widetilde{V}}}(W, \sqrt{n}\epsilon/4)\right)}. \tag{42}$$

In particular, we choose π to the uniform prior over $\overline{\mathcal{V}}$:

$$\pi_{\bar{\mathcal{V}}} = \text{Unif}(B_2(1.58R) \cap \mathcal{V}).$$

Then we prove that $B_{\varrho_{\Sigma}^{\bar{\mathcal{V}}}}(W,\sqrt{n}\epsilon/4)\subseteq \bar{\mathcal{V}}\cap B_2(1.58R)$. This is because: 1) for every $W'\in B_{\varrho_{\Sigma}^{\mathcal{V}}}(\mathcal{P}_{\mathcal{V}}(W),\sqrt{n}\epsilon/4)$, (41) suggests $\|W'-\mathcal{P}_{\mathcal{V}}(W)\|_2^2\leq \frac{1}{3}R^2$, and 2) for very $W\in\mathcal{W}$, we have $\|\mathcal{P}_{\mathcal{V}}(W)\|_2\leq \|W\|_2^2\leq R^2$. Combining this and the above inequality we have

$$\|W'\|_2 \le \|W' - \mathcal{P}_{\mathcal{V}}(W)\|_2 + \|\mathcal{P}_{\mathcal{V}}(W)\|_2 \le (\sqrt{1/3} + 1)R < 1.58R.$$

This proves that $B_{\varrho \Sigma}(\mathcal{P}_{\mathcal{V}}(W), \sqrt{n}\varepsilon/4) \subseteq \mathcal{V} \cap B_2(1.58R)$, so we have

$$\log \frac{1}{\pi_{\bar{\mathcal{V}}}(B_{\varrho_{\bar{\mathcal{V}}}^{\bar{\mathcal{V}}}}(W,\sqrt{n}\epsilon/4))} = \frac{\operatorname{Vol}(\bar{\mathcal{V}} \cap B_2(1.58R))}{\operatorname{Vol}(B_{\varrho_{\bar{\mathcal{V}}}^{\bar{\mathcal{V}}}}(W,\sqrt{n}\epsilon/4))}$$
(43)

Applying Lemma 17 about the volume of ellipsoid, with dimension $r_{\rm eff}$, eigenvalues $\{\lambda_k(\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}})\}_{k=1}^{r_{\rm eff}}$ and radius $r=\sqrt{n}\varepsilon/4$:

$$\begin{aligned} \operatorname{Vol}(B_{\varrho_{\Sigma}^{\mathcal{V}}}(\mathcal{P}_{\mathcal{V}}(W), \sqrt{n/2}\varepsilon)) = & \frac{\pi^{r_{\text{eff}}/2}}{\Gamma(r_{\text{eff}}/2+1)} \Big(\sqrt{n} \cdot \varepsilon/4\Big)^{r_{\text{eff}}} \Big(\Pi_{k=1}^{r_{\text{eff}}} \lambda_{k} (\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}})\Big)^{-1/2} \\ \geq & \frac{\pi^{r_{\text{eff}}/2}}{\Gamma(r_{\text{eff}}/2+1)} \Big(\sqrt{n} \cdot \varepsilon/4\Big)^{r_{\text{eff}}} \Big(\Pi_{k=1}^{r_{\text{eff}}} \lambda_{k}\Big)^{-1/2}, \end{aligned}$$

where the inequality uses Lemma 18 to upper bound $\lambda_k(\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}})$ by λ_k . The standard formula for the volume of unit ball, Lemma 17, suggests that the volume of $\mathcal{V} \cap B_2(2R)$ is

$$Vol(\mathcal{V} \cap B_2(1.58R)) = \frac{\pi^{r_{\text{eff}}/2}}{\Gamma(r_{\text{eff}}/2+1)} (1.58R)^{r_{\text{eff}}}.$$

Hence, applying (42) (43) and combining it with the two above volume bounds, we have

$$\begin{split} \log \frac{1}{\pi \left(B_{\varrho_{\Sigma}}(W,\sqrt{n}\varepsilon)\right)} &\leq \log \frac{1}{\pi_{\bar{\mathcal{V}}}(B_{\varrho_{\Sigma}^{\bar{\mathcal{V}}}}(\mathcal{P}_{\mathcal{V}}(W),\sqrt{n}\varepsilon/4))} = \log \frac{\operatorname{Vol}(\mathcal{V} \cap B_{2}(2R))}{\operatorname{Vol}(B_{\varrho_{\Sigma}^{\mathcal{V}}}(\mathcal{P}_{\mathcal{V}}(W),\sqrt{n}\varepsilon/4))} \\ &\leq \frac{1}{2}\log \frac{(1.58R)^{2r_{\mathrm{eff}}} \Pi_{k=1}^{r_{\mathrm{eff}}} \lambda_{k}}{(\sqrt{n}\varepsilon/4)^{2r_{\mathrm{eff}}}} &\leq \frac{1}{2}\sum_{k=1}^{r_{\mathrm{eff}}} \log \frac{40R^{2}\lambda_{k}}{n\varepsilon^{2}} \\ &= d_{\mathrm{eff}}(\Sigma,\sqrt{5}R,\epsilon). \end{split}$$

Finally, since the prior construction $\pi_{\bar{\mathcal{V}}} = \mathrm{Unif}(B_2(1.58R) \cap \bar{\mathcal{V}})$ only depends on $\bar{\mathcal{V}}$ rather than Σ and ϵ , we have that uniformly over all $(W, \epsilon) \in \mathcal{W} \times [0, \infty)$ such that $\bar{\mathcal{V}}$ approximates $\mathcal{V}_{\mathrm{eff}}(G(W), R, \epsilon)$ to the precision (37),

$$\log \frac{1}{\pi_{\bar{\mathcal{V}}}(B_{\varrho_{\Sigma}}(W, \sqrt{n\varepsilon}))} \le d_{\text{eff}}(G(W), \sqrt{5}R, \epsilon).$$

Proof of Lemma 18: The Courant–Fischer–Weyl max-min characterization Wikipedia contributors (b) states that for any Hermitian (i.e. symmetric for real matrices studying here) matrix,

$$\lambda_k(\Sigma) = \max_{\substack{S \subseteq \mathbb{R}^p \\ \dim S = k}} \min_{\substack{W \in S \\ W \neq 0}} \frac{W^{\top} \Sigma W}{\|W\|^2},$$

and we have that for any r-dimensional sub-space $\bar{\mathcal{V}}$,

$$\lambda_k(\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}}) = \max_{\substack{\mathcal{S} \subseteq \bar{\mathcal{V}} \\ \dim \mathcal{S} = k}} \min_{\substack{W \in \mathcal{S} \\ W \neq 0}} \frac{W^{\top} M W}{\|W\|^2},$$

so we have $\lambda_k(\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}}) \leq \lambda_k$ for $k = 1, 2, \cdots, r$.

Moreover, by the elementary property of eigendecomposition we have $\lambda_k = \lambda_k(\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}})$, and by the Courant–Fischer–Weyl max-min characterization we know that,

$$\lambda_{k}(\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}}) = \max_{\substack{\mathcal{S}\subseteq\mathbb{R}^{p}\\\dim\mathcal{S}=k}} \min_{\substack{W\in\mathcal{S}\\W\neq 0}} \frac{W^{\top}(\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}})W}{\|W\|^{2}}$$

$$\leq \max_{\substack{\mathcal{S}\subseteq\mathbb{R}^{p}\\\dim\mathcal{S}=k}} \min_{\substack{W\in\mathcal{S}\\W\neq 0}} \frac{W^{\top}(2\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}})W + \|\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}} - 2\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}}\|_{op}\|W\|^{2}}{\|W\|^{2}}$$

$$= 2\lambda_{k}(\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}}) + \|\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}} - 2\mathcal{P}_{\bar{\mathcal{V}}}^{\top}\Sigma\mathcal{P}_{\bar{\mathcal{V}}}\|_{op}$$

$$\leq 2\lambda_{k}(\mathcal{P}_{\mathcal{V}}^{\top}\Sigma\mathcal{P}_{\mathcal{V}}) + 2\|(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})^{\top}\Sigma(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{op},$$

where the last inequality is due to (39). Therefore we have

$$\lambda_k(\mathcal{P}_{\bar{\mathcal{V}}}^{\top} \Sigma \mathcal{P}_{\bar{\mathcal{V}}}) \ge \lambda_k/2 - \|\Sigma^{\frac{1}{2}} (\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{op}}^2.$$

C.4 PROOF FOR NN-TYPE MANIFOLD (THEOREM 3)

For a NN-type manifold, by its definition in Theorem 3 we have the structure that the metric tensor G(W) is

$$G(W) = blockdiag(A_1(W) \otimes I_{d_1}; \cdots; A_l(W) \otimes I_{d_l}; \cdots; A_L(W) \otimes I_{d_L}),$$

where each A_l is a $d_{l-1} \times d_{l-1}$ PSD matrix. We firstly state straightforward decomposition properties for such manifold, and then prove Theorem 3.

C.4.1 DECOMPOSITION PROPERTIES OF NN-TYPE MANIFOLD

The structured G(W) in NN-type Manifold has decomposition properties described by the next lemma.

Definition 3 (Metric Tensor for NN-type Manifold) A manifold (W, G) is called "NN-type" if $W \subseteq \mathbb{R}^{\sum_{l=1}^{L} d_{l-1}d_{l}}$ and its metric tensor G(W) is in the form

$$G(W) = blockdiag(A_1(W) \otimes I_{d_1}, \cdots, A_l(W) \otimes I_{d_l}, \cdots, A_L(W) \otimes I_{d_L})$$

where $A_l(W) \in \mathbb{R}^{d_{l-1} \times d_{l-1}}$.

Lemma 19 (Decomposition Properties of NN-type Manifold) Consider a NN-type manifold (W, G) defined in Definition 3. For every $W \in W$, we have the following decomposition properties: First, the effective rank and dimension decompose to

$$r_{\text{eff}}(G(W), R, \epsilon) = \sum_{l=1}^{L} d_l \cdot r_{\text{eff}}(A_l(W), R, \epsilon);$$
$$d_{\text{eff}}(G(W), R, \epsilon) = \sum_{l=1}^{L} d_l \cdot d_{\text{eff}}(A_l(W), R, \epsilon).$$

Second, denote $V_{\text{eff}}(A_l(W), R, \epsilon)$ the effective sub-space of $A_l(W)$. Then the effective sub-space of G(W) is

$$\mathcal{V}_{\text{eff}}(G(W), R, \epsilon) = \mathcal{V}_{\text{eff}}(A_1(W), R, \epsilon)^{d_1} \times \cdots \times \mathcal{V}_{\text{eff}}(A_L(W), R, \epsilon)^{d_L}.$$

Proof of Lemma 19. It is straightforward to see that, first, the effective rank of the fixed matrix G(W) is

$$r_{\text{eff}}(G(W), R, \epsilon)$$

$$= \max\{k : 2\lambda_k(G(W))R^2 \ge n\epsilon^2\}$$

$$= \sum_{l=1}^L \max\{k : 2\lambda_k(A_l(W) \otimes I_{d_l})R^2 \ge n\epsilon^2\}$$

$$= \sum_{l=1}^L d_l \max\{k : 2\lambda_k(A_l(W))R^2 \ge n\epsilon^2\}$$

$$= \sum_{l=1}^L d_l \cdot r_{\text{eff}}(A_l(W), R, \epsilon);$$

and the effective dimension of the fixed matrix G(W) is

$$\begin{split} & d_{\text{eff}}(G(W), R, \epsilon) \\ &= \frac{1}{2} \sum_{k=1}^{r_{\text{eff}}(G(W), R, \epsilon)} \log \left(\frac{8R^2 \lambda_k(G(W))}{n\epsilon^2} \right) \\ &= \sum_{l=1}^{L} \frac{1}{2} \sum_{k=1}^{r_{\text{eff}}(A_l(W) \otimes I_{d_l}, R, \epsilon)} \log \left(\frac{8R^2 \lambda_k(A_l(W) \otimes I_{d_l})}{n\epsilon^2} \right) \\ &= \sum_{l=1}^{L} d_l \cdot \frac{1}{2} \sum_{k=1}^{r_{\text{eff}}(A_l(W), R, \epsilon)} \log \left(\frac{8R^2 \lambda_k(A_l(W))}{n\epsilon^2} \right) \\ &= \sum_{l=1}^{L} d_l \cdot r_{\text{eff}}(A_l(W), R, \epsilon). \end{split}$$

Second, as the effective sub-space of the matrix tensor product $A_l(W) \otimes I_{d_l}$ is sub-space tensor product $\mathcal{V}_{\text{eff}}(A_l(W), R, \epsilon)^{d_l}$, the effective sub-space for NN-type metric tensor $G(W) = \text{blockdiag}(\cdots; A_l(W) \otimes I_{d_l}; \otimes)$ is

$$V_{\rm eff}(G(W),R,\epsilon) := \mathcal{V}_{\rm eff}(A_1(W),R,\epsilon)^{d_1} \times \cdots \times \mathcal{V}_{\rm eff}(A_L(W),R,\epsilon)^{d_L}.$$

We firstly prove the following result, which is almost Theorem 3, with the only difference being that the radius in the effective dimension depends on the global radius R rather than the pointwise Frobenious norm $\|W\|$. Extending this result to Theorem 3 can be achieved via a simple application of the "uniform pointwise convergence" principle Xu & Zeevi (2025) illustrated in Lemma 8.

Lemma 20 (Riemannian Dimension for NN-type Manifold—Global Radius Version)

Consider a NN-type manifold (W,G) and assume $\sup_{\mathcal{W}} ||W||_{\mathsf{F}} \leq R$. Then (W,G) admits Riemannian Dimension

$$d_R(W,\epsilon) = \sum_{l=1}^L \bigg(\underbrace{d_l \cdot d_{\mathrm{eff}}(A_l(W), CR, \epsilon)}_{\text{"must pay" cost at each }W} + \underbrace{d_{l-1} \cdot d_{\mathrm{eff}}(A_l(W), CR, \epsilon)}_{\text{covering cost of Grassmannian}} + \underbrace{\log(d_{l-1})}_{\text{covering cost of }r_{\mathrm{eff}} \in [d_{l-1}]} \bigg),$$

where C > 0 is an absolute constant.

C.4.2 PROOF OF THEOREM 3

Proof of Lemma 20: The proof has two key steps: 1. Hierarchical covering argument, and 2. Bound covering Cost of the Grassmannian. A crucial lemma about the ellipsoidal covering of the Grassmannian, which is new even in the pure mathematics context, is deferred to Section D.

Step 1: Hierarchical Covering. As explained the main paper, the major difficulty is that the prior measure $\pi_{\mathcal{V}}$ it constructed, is defined over the effective sub-space \mathcal{V} , which itself encodes information of the point $W \in \mathcal{W}$ and $\epsilon > 0$. The goal of our proof is to construct a "universal" prior π that does not depend on \mathcal{V} . This is achieved via a hierarchical covering argument (10), which we make rigorous below.

The key idea of hierarchical covering is as follows: Firstly, for all $W \in \mathcal{W}$, we search for subspace $\bar{\mathcal{V}}$ that approximates the true effective sub-space (top- $r_{\rm eff}$ eigenspace) $\mathcal{V}_{\rm eff}(G(W), R, \epsilon)$ to the precision required by (37):

$$||G(W)^{\frac{1}{2}}(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})||_{\text{op}} \le \frac{\sqrt{n}\epsilon}{4R},\tag{44}$$

where $G(W)^{\frac{1}{2}}$ is the unique square root of PSD matrix G(W) (see, e.g, wikipedia Wikipedia contributors (c)). Then by Lemma 2 (Pointwise Dimension Bound for Non-Linear Manifold with Approximate Effective Sub-Space), for every $(W,\epsilon) \in \mathcal{W} \times [0,\infty)$ such that $\bar{\mathcal{V}}$ approximates $\mathcal{V}_{\mathrm{eff}}(G(W),R,\epsilon)$ to the precision (44), the prior $\pi_{\bar{\mathcal{V}}}=\mathrm{Unif}(B_2(1.58R)\cap\bar{\mathcal{V}})$ satisfies

$$\log \frac{1}{\pi_{\bar{V}}(B_{\varrho_{\Sigma}}(w,\sqrt{n}\varepsilon))} \leq d_{\text{eff}}(G(W),\sqrt{5}R,\epsilon) = \sum_{l=1}^{L} d_{l} \cdot d_{\text{eff}}(A_{l}(W),\sqrt{5}R,\epsilon), \quad (45)$$

where the first inequality is by Lemma 2; the first equality is by definition (8) of effective dimension; and the last equality is by the decomposition property of NN-type manifold (Lemma 19).

Secondly, we put a prior μ over all possible sub-spaces $\mathcal V$ and construct the "universal" prior

$$\pi(W) = \sum_{\mathcal{V}} \mu(\mathcal{V}) \times \pi_{\mathcal{V}}(W), \tag{46}$$

which implies that uniformly over all $W \in \mathcal{W}$,

$$\log \frac{1}{\pi(B_{\varrho_{G(W)}}(W, \sqrt{n}\epsilon))}$$

$$= \log \frac{1}{\sum_{\mathcal{V}} \mu(\mathcal{V}) \pi_{V}(B_{\varrho_{G(W)}}(W, \sqrt{n}\epsilon))}$$

$$\leq \log \frac{1}{\mu(\bar{\mathcal{V}} : \bar{\mathcal{V}} \text{ satisfies (44)}) \inf_{\bar{\mathcal{V}} \text{ satisfies (44)}} \pi_{\bar{\mathcal{V}}}(B_{\varrho_{G(W)}}(W, \sqrt{n}\epsilon))}$$

$$\leq \log \frac{1}{\mu(\bar{\mathcal{V}} : \bar{\mathcal{V}} \text{ satisfies (44)})} + \sum_{l=1}^{L} d_{l} \cdot d_{\text{eff}}(A_{l}(W), \sqrt{5}R, \epsilon), \tag{47}$$

where the first equality is by definition (46) of the "universal" prior π ; the first inequality is straightforward; and the last inequality is by (45), the result of the "must pay" part in the hierarchical covering.

The above hierarchical covering argument successfully gives a valid Riemannian Dimension, with the cost of the additional covering cost given by the sub-space prior μ . This explains our basic proof idea. The remaining proof executes this basic proof idea.

Step 2: Bounding Covering Cost of the Grassmannian. Section D provides a systematical study to the ellipsoidal metric entropy of Grassmannian manifold, which we detail the conclusion below.

Define

$$\operatorname{Gr}(d,r) := \{r \text{-dimensional linear subspaces of } \mathbb{R}^d \}$$

as the Grassmann manifold.

Given a $d \times d$ PSD Σ , define the anisometric projection metric between two sub-spaces by (labeled as Definition 4 in Section D)

$$\varrho_{\text{proj},\Sigma}(\mathcal{V},\bar{\mathcal{V}}) = \|\Sigma^{\frac{1}{2}}(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{op}},\tag{48}$$

where $\Sigma^{\frac{1}{2}}$ is the square root of the PSD matrix Σ (see, e.g., wikipedia Wikipedia contributors (c)).

Lemma 3 states that, given a Grassmannian Gr(d,r), for uniform prior $\mu = Unif(Gr(d,r))$, we have that for every $\mathcal{V} \in Gr(d,r)$, every $\epsilon > 0$ and PSD matrix $\Sigma \in \mathbb{R}^{d \times d}$, we have the pointwise dimension bound

$$\log \frac{1}{\mu(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V},\epsilon))} \le \frac{1}{2}(d-r)\sum_{k=1}^{r}\log \frac{C\max\{\lambda_{k},\epsilon^{2}/42\}}{\epsilon^{2}} + r\sum_{k=1}^{d-r}\log \frac{C\max\{\lambda_{k},\epsilon^{2}/42\}}{\epsilon^{2}},\tag{49}$$

where $C_1 > 0$ is an absolute constant. We will use the result (49) and (47) to prove Theorem 3.

For a particular layer l, $d_{l-1} \times d_{l-1}$ PSD matrix $A_l(W)$, and a fixed rank r_l denote $\operatorname{Gr}(d_{l-1}, r_l)$ as a Grassmannian (the collection of all r_l -dimensional in $\mathbb{R}^{d_{l-1}}$). By (49) we have that there exists a prior μ_l over $\operatorname{Gr}(d_{l-1}, r_l)$ such that for every $W \in \mathcal{W}$ such that $r_{\operatorname{eff}}(A_l(W), R, \epsilon) = r_l$, and $\epsilon_1 \leq \lambda_{r_l}(A_l(W))$, we have

$$\log \frac{1}{\mu_l(\bar{\mathcal{V}}: \varrho_{\text{proj}, A_l(W)}(\mathcal{V}_{\text{eff}}(A_l(W), R, \epsilon), \bar{\mathcal{V}}) \le \epsilon_1)} \le \frac{1}{2} d_{l-1} \sum_{k=1}^{r_l} \log \frac{C_1 \lambda_k(A_l(W))}{\epsilon_1^2}. \tag{50}$$

By the sub-space decomposition property in Lemma 19, we have that for $\bar{\mathcal{V}} = (\cdots, \underline{\bar{\mathcal{V}}_l, \cdots, \bar{\mathcal{V}}_l}, \cdots)$, repeat d_l times

$$\varrho_{\text{proj},G(W)}(\mathcal{V}_{\text{eff}}(G(W), R, \epsilon), \bar{\mathcal{V}})
= \varrho_{\text{proj},G(W)}(\Pi_{l=1}^{L} \mathcal{V}_{\text{eff}}(A_{l}(W), R, \epsilon)^{d_{l}}, \Pi_{l=1}^{L} \bar{\mathcal{V}}_{l}^{d_{l}})
= \max_{l} \varrho_{\text{proj},A_{l}(W)}(\mathcal{V}_{\text{eff}}(A_{l}(W), R, \epsilon), \bar{\mathcal{V}}),$$
(51)

where the first equality is by Lemma 19, and the second equality is by the properties of the spectral norm: $\|\text{blockdiag}(A, B)\|_{\text{op}} = \max\{\|A\|_{\text{op}}, \|B\|_{\text{op}}\}$ and $\|A \otimes I_d\|_{\text{op}} = \|A\|_{\text{op}}$.

Taking $\epsilon_1 = \frac{\sqrt{n}\epsilon}{4R}$, we see this choice satisfies the required condition to establish (50). Then for all layers $l=1,\cdots,L$, given a fixed $\{r_1,\cdots,r_L\}$, by (50), we have that there exists a prior

$$\mu_{\{r_l\}_{l=1}^L} = \mu_1^{d_1} \times \dots \times \mu_L^{d_L} = \prod_{l=1}^L (\underbrace{\mu_l \times \dots \times \mu_l}_{d_l \text{ times}})$$
 (52)

over the product Grassmannian $\operatorname{Gr}(d_0,r_1)^{d_1}\times\cdots\times\operatorname{Gr}(d_{L-1},r_L)^{d_L}$ such that univerformly over all $W\in\mathcal{W}$ such that $r_{\operatorname{eff}}(A_l(W),R,\epsilon)=r_l,\,\forall l\in[L]$ (here [L] is the notation of $\{1,2,\cdots,L\}$), the "Grassmannian covering cost" term in (47) is bounded by

$$\log \frac{1}{\mu(\bar{\mathcal{V}}: \bar{\mathcal{V}} \text{ satisfies (44)})}$$

$$= \log \frac{1}{\mu_{\{r_l\}_{l=1}^L}(\bar{\mathcal{V}}: \varrho_{\text{proj},G(W)}(\mathcal{V}_{\text{eff}}(G(W), R, \epsilon), \bar{\mathcal{V}}) \leq \frac{\sqrt{n}\epsilon}{4R} = \epsilon_1)}$$

$$\leq \log \frac{1}{\mu_{\{r_l\}_{l=1}^L}((\cdots, \bar{\mathcal{V}}_l, \cdots, \bar{\mathcal{V}}_l, \cdots) : \varrho_{\text{proj},A_l(W)}(\mathcal{V}_{\text{eff}}(A_l(W), R, \epsilon), \bar{\mathcal{V}}_l) \leq \epsilon_1, \quad \forall l \in [L])}$$

$$= \sum_{l=1}^L \log \frac{1}{\mu_{\{r_l\}_{l=1}^L}((\cdots, \bar{\mathcal{V}}_l, \cdots, \bar{\mathcal{V}}_l, \cdots) : \varrho_{\text{proj},A_l(W)}(\mathcal{V}_{\text{eff}}(A_l(W), R, \epsilon), \bar{\mathcal{V}}_l) \leq \epsilon_1)}$$

$$\leq \sum_{l=1}^L d_{l-1} \sum_{k=1}^{r_l} \log \frac{C_1 \lambda_k(A_l(W))}{\epsilon_1^2}$$

$$\leq \sum_{l=1}^L d_{l-1} d_{\text{eff}}(A_l(W), \sqrt{2C_1}R, \epsilon), \tag{53}$$

where the first inequality is by restricting $\mathcal V$ to the from $\Pi_{l=1}^L \bar{\mathcal V}_l^{d_l}$ and using (51); the second equality is by the choice of he product prior (52); the last inequality is by the layer-wise covering bound (50); and the second inequality is by 1) $d_{l-1}-r_l \leq d_{l-1}$ and $C_1\lambda_k(A_l(W)) \geq \lambda_{r_l} \geq \epsilon_1$ (as stated in (49)), 2) the choice $\epsilon_1 = \sqrt{n}\epsilon/(4R)$, and 3) definition (8) of effective dimension.

Note that (53) is uniformly over all $W \in \mathcal{W}$ such that $r_{\text{eff}}(A_l(W), R, \epsilon) = r_l, \forall l \in [L]$, not uniformly over all $W \in \mathcal{W}$. We would like to extend (53) to all $W \in \mathcal{W}$ over uniform prior over possible integer values of r_l . Now assign uniform prior over $[d_{l-1}] = \{1, \cdots, d_{l-1}\}$ for r_l , we obtain the "universal" prior π (as we have pursued in in our hierarchical covering argument (46)) defined by

$$\mu(\mathcal{V}) = \prod_{l=1}^{L} \underbrace{\mathrm{Unif}([d_{l-1}])}_{\mathrm{prior \ of} \ r_{l}} \times \underbrace{\mu_{\{r_{k}\}_{k=1}^{L}}}_{\mathrm{prior \ over \ product \ Grassmannian \ in \ (52)}}_{\mathrm{prior \ over \ sub-spaces \ defined \ above}} \times \underbrace{\mathrm{Unif}(B_{2}(0, 1.58R) \cap \bar{V})}_{\mathrm{uniform \ prior \ constrained \ in \ sub-space}}. \tag{54}$$

Then we have that uniformly over all $W \in \mathcal{W}$,

$$\begin{split} &\log \frac{1}{\pi(B_{\varrho_{G(W)}}(W,\sqrt{n}\epsilon))} \\ &\leq \log \frac{1}{\mu(\bar{\mathcal{V}}:\bar{\mathcal{V}} \text{ satisfies (44)})} + \sum_{l=1}^{L} d_l \cdot d_{\text{eff}}(A_l(W),\sqrt{5}R,\epsilon)) \\ &\leq \sum_{l=1}^{L} \log d_{l-1} + \log \frac{1}{\mu_{\{r_k\}_{k=1}^{L}}(\bar{\mathcal{V}}:\bar{\mathcal{V}} \text{ satisfies (44)})} + \sum_{l=1}^{L} d_l \cdot d_{\text{eff}}(A_l(W),\sqrt{5}R,\epsilon)) \\ &\leq \sum_{l=1}^{L} \log d_{l-1} + \sum_{l=1}^{L} d_{l-1} \cdot d_{\text{eff}}(A_l(W),\sqrt{2C_1}R,\epsilon) + + \sum_{l=1}^{L} d_l \cdot d_{\text{eff}}(A_l(W),\sqrt{5}R,\epsilon)), \end{split}$$

where $C_1 > 1$ is an absolute constant. Here the first inequality is by the hierarchical covering argument (47); the second inequality is by the prior construction (54); and the third inequality is by the Grassmannian covering bound (53) for fixed $\{r_k\}_{k=1}^L$. This shows that the NN-type manifold (\mathcal{W}, G) admits Riemannian Dimension

$$\log \frac{1}{\pi(B_{\varrho_{G(W)}}(W, \sqrt{n}\epsilon))} \le \sum_{l=1}^{L} (d_l + d_{l-1}) \cdot d_{\text{eff}}(A_l(W), CR, \epsilon) + \log d_{l-1}),$$

where C is a positive absolute constant. This finishes the proof of Lemma 20 with R in effective dimension being a global upper bound of $||W||_{\mathbf{F}}$.

Proof of Theorem 3: Motivated by the "uniform pointwise convergence" principle (proposed in Xu & Zeevi (2025) and illustrated in Lemma 8), we apply a peeling argument to adapt the Riemannian Dimension to $\|W\|_{\mathbf{F}}$. Given any $R_0 \in (0,R]$, we take $R_k = 2^k R_0$ for $k = 0,1, \cdots \log_2 \lceil Rn \rceil$. Taking a uniform prior on these R_k , and set

$$\tilde{\pi} = \underbrace{\mathrm{Unif}(\{R_0, \cdots, 2^{\log_2 \lceil R/R_0 \rceil} R_0\})}_{\text{prior over upper bound } \tilde{R} \text{ of } \|W\|_{\mathbf{F}}} \times \underbrace{\pi_{\tilde{R}}}_{\text{prior defined via (54)}},$$

where π_R is the prior defined via (54) in the proof of Lemma 20. Then for every $W \in \mathcal{W}$ where $\|W\|_{\mathbf{F}} > R_0$, denote k(W) to be the integer such that $2^{k(W)}R_0 < \|W\|_{\mathbf{F}} \le 2^{k(W)+1}R_0$, then

$$\begin{split} & \log \frac{1}{\tilde{\pi}(B_{\varrho_{G(W)}}(W,\sqrt{n}\epsilon))} \\ & \leq \underbrace{\log \log_2 \lceil R/R_0 \rceil}_{\text{density of } 2^{k(W)+1}R_0} + \underbrace{\log \frac{1}{\pi_{2^{k(W)+1}R_0}(B_{\varrho_{G(W)}}(W,\sqrt{n}\epsilon))}}_{\pi \text{ is constructed via (54), with global radius taken to be } 2^{k(W)+1}R_0} \\ & \leq \log \log_2 \lceil Rn \rceil + \sum_{l=1}^L (d_l + 2d_{l-1}) \cdot d_{\text{eff}}(A_l(W), C_1 2^{k_0+1}R_0, \epsilon) \log d_{l-1}) \\ & \leq \log \log_2 \lceil Rn \rceil + \sum_{l=1}^L (d_l + 2d_{l-1}) \cdot d_{\text{eff}}(A_l(W), C_1 2^{k_0+1}R_0, \epsilon) + \log d_{l-1}) \end{split}$$

where the first inequality is due to the product construction of $\bar{\pi}$; the second inequality is due to Lemma 20, with $C_1>0$ being an absolute constant; and the last inequality uses the fact $\|W\|_{\mathbf{F}}\leq 2^{k_0+1}R_0\leq 2\|W\|_{\mathbf{F}}$, with $C_1>0$.

The above bound assumes $||W||_{\mathsf{F}} > R_0$. When $||W|| \le R_0$, we directly apply Lemma 20 and obtain

$$\log \frac{1}{\tilde{\pi}(B_{\varrho_{G(W)}}(W, \sqrt{n}\epsilon))} \le \underbrace{\log \log_2 \lceil R/R_0 \rceil}_{\text{density of } R_0} + \underbrace{\log \frac{1}{\pi_{R_0}(B_{\rho_{G(W)}}(W, \sqrt{n}\epsilon))}}_{\pi \text{ is constructed via (54), with global radius taken to be } R_0$$

$$\leq \log \log_2 \lceil Rn \rceil + \sum_{l=1}^{L} (d_l + 2d_{l-1}) \cdot d_{\text{eff}}(A_l(W), C_1 \cdot R_0, \epsilon) + \log d_{l-1}).$$

Combining the two cases discussed above, we conclude that the NN-type manifold (\mathcal{W},G) admits Riemmanin Dimension

$$\begin{split} & d_{R}(W, \epsilon) \\ &= \sum_{l=1}^{L} (d_{l} + 2d_{l-1}) \cdot d_{\text{eff}}(A_{l}(W), C \max\{\|W\|_{\mathbf{F}}, R_{0}\}) + \log(d_{l-1} \log_{2} \lceil R/R_{0} \rceil), \end{split}$$

where $C=2C_1$ is a positive absolute constant. Taking $R_0=R/2^n$ proves Theorem 3.

D ELLIPSOIDAL COVERING OF THE GRASSMANNIAN

The central goal of this section is to prove the following result on the ellipsoidal metric entropy of the Grassmannian manifold. The definition for Gr (Grassmannian manifold), St (Stiefel parameterization manifold) are temporarily deferred to Section D.1.

Definition 4 (Ellipsoidal Projection Metric) For two sub-spaces $V, \bar{V} \in Gr(d, r)$, and a positive semidefinite matrix Σ , define the ellipsoidal projection metric $\varrho_{\text{proj},\Sigma}$ by

$$\varrho_{\mathrm{proj},\Sigma}(\mathcal{V},\bar{\mathcal{V}}) = \|\Sigma^{\frac{1}{2}}(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\mathrm{op}}$$

where $\mathcal{P}_{\mathcal{V}}$ and $\mathcal{P}_{\bar{\mathcal{V}}}$ are orthogonal projectors to sub-space \mathcal{V} and $\bar{\mathcal{V}}$, respectively.

We view orthogonal projectors as matrices (see the definition via the Stiefel parameterization in (58)), consistent with the earlier operator notation characterized by ℓ_2 -distance in Lemma 14. In the isotropic case $\Sigma = I_d$, the ellipsoidal projection metric reduces to the standard isotropic projection metric

$$\varrho_{\text{proj}}(\mathcal{V}, \bar{\mathcal{V}}) = \left\| \mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}} \right\|_{\text{op}}.$$

We now state our main result in this section (Lemma 3 in the main paper).

Consider the Grassmannian $\operatorname{Gr}(d,r)$. For uniform prior $\mu=\operatorname{Unif}(\operatorname{Gr}(d,r))$, we have that for every $\mathcal{V}\in\operatorname{Gr}(d,r)$, every $\epsilon>0$ and every PSD matrix Σ , we have the pointwise dimension bound

$$\log \frac{1}{\mu(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V},\epsilon))} \leq \frac{1}{2}(d-r)\sum_{k=1}^r \log \frac{C \max\{\lambda_k,\epsilon^2/42\}}{\epsilon^2} + r\sum_{k=1}^{d-r} \log \frac{C \max\{\lambda_k,\epsilon^2/42\}}{\epsilon^2},$$

where C > 0 is an absolute constant.

Recall that the traditional covering number bound for the Grassmannian manifold states that

$$\left(\frac{c}{\epsilon}\right)^{r(d-r)} \le \mathcal{N}(\operatorname{Gr}(d,r), \varrho_{\operatorname{proj}}, \epsilon) \le \left(\frac{C}{\epsilon}\right)^{r(d-r)}. \tag{55}$$

Here $N(\mathcal{F}, \varrho, \varepsilon)$ is the standard covering number— the smallest size of an ε -net that covers \mathcal{F} under the metric ϱ ; see Definition 4.2.2 in Vershynin (2009) for details. In comparison, Lemma 3 is much more challenging than proving classical isotropic covering number bounds (55) because

• 1) we consider ellipsoidal metric;

 • 2) we require the prior μ to be independent with $\Sigma(\mathcal{V})$ and ϵ .

We need to firstly understand how such classical results are proved, and then proceed to generalized them. This suggests that deep mathematical insights are necessary for the purpose to study neural networks generalization, as we will introduce below.

From Pure Mathematics to Machine Learning Language. Understanding the classical proof for the Grassmannian and generalizing them to prove Lemma 3 necessitate the a deep dive in to the geometry and algebra of sub-spaces and Grassmannians. In fact, traditional treatments to study Grassmannian manifold often invoke advanced machinery—ranging from differential geometry Bendokat et al. (2024) and Lie-group theory Szarek (1997) to algebraic geometry Devriendt et al. (2024), and the seminal covering number proof Szarek (1997) is particularly stated in Lie-algebra and differential-geometry language.

Motivated by the subsequent covering number proof Pajor (1998) that uses relatively more elementary language, we give an exposition that is elementary and entirely self-contained, relying only on matrix-analysis and learning-theoretic techniques familiar from machine-learning. In particular, every "advanced" fact is derived by elementary means (careful matrix parameterizations and basic spectral arguments) while preserving high-level intuition. We hope that this versatile framework—and our novel contributions (e.g., Definition 4 and Lemma 3), which are new even in a puremathematics setting—will establish subspaces, the Grassmannian, and their underlying algebraic structures as powerful tools for future machine-learning applications.

D.1 GRASSMANNIAN MANIFOLD, STIEFEL PARAMETERIZATION, AND ORTHOGONAL GROUPS

Fix integers $r \leq d$. Define

$$Gr(d, r) := \{r - \text{dimensional linear subspaces of } \mathbb{R}^d \}$$

as the Grassmann manifold. Write

$$\mathrm{St}(d,r) \; := \; \left\{ \, V \in \mathbb{R}^{d \times r} \; : \; V^\top V = I_r \, \right\}$$

for the *Stiefel manifold* of r orthonormal columns in \mathbb{R}^d . St(n,m) is a convenient *parameterization* of that class Gr(d,r).

If for sub-space $V \in Gr(d,r)$ and matrix $V \in St(d,r)$ we have $V = \operatorname{span}(V)$, then we say V is a *parameterization matrix* of V. Though such parameterization is not unique, the associated orthogonal projector and projection metric are both unique. Moreover, the anisometric projection we define in Definition 4 is also unique. We will prove these shortly.

Write

$$O(r) := \{ Q \in \mathbb{R}^{r \times r} : Q^{\top}Q = QQ^{\top} = I_r \}$$

to be the *orthogonal group*. Optionally, we also state that (in the real setting)

$$\operatorname{Gr}(d,r) \cong O(d) / (O(r) \times O(d-r)) \cong \operatorname{Gr}(d,d-r),$$
 (56)

where "/" denotes the *quotient* and " \cong " denotes a canonical *isomorphism* (indeed, a *diffeomorphism* of smooth manifolds or a *homeomorphism* of topological manifolds; see, e.g., Chapter 1.5 in Awodey (2010)). Moreover, Gr(d, r) can be regarded as a standard *algebraic variety* Devriendt et al. (2024). We do not aim to explain these notions in detail, but merely note that:

- 1. The geometric properties of Gr(d, r) coincide with those of Gr(d, d r) under this isomorphism (geometric equivalence).
- 2. The number of degrees of freedom of Gr(d, r) is

$$\underbrace{\frac{d(d-1)}{2}}_{\dim O(d)} - \underbrace{\frac{r(r-1)}{2}}_{\dim O(r)} - \underbrace{\frac{(d-r)(d-r-1)}{2}}_{\dim O(d-r)} = r(d-r), \tag{57}$$

which also appears as the dimension factor in the precise covering-number bounds (55).

We now define the orthogonal projector and the projection metric on the Grassmannian manifold.

Definition of Orthogonal Projector. For $V \in St(d, r)$ and its column-space $\mathcal{V} = \operatorname{span}(V)$, define the rank-r orthogonal projector¹

$$P_{\mathcal{V}} := VV^{\top} \in \mathbb{R}^{d \times d}. \tag{58}$$

Then $P_{\mathcal{V}}$ depends only on the sub-space \mathcal{V} . Indeed, if $Q \in O(r)$ then $(VQ)(VQ)^{\top} = VQQ^{\top}V^{\top} = VV^{\top}$, so V and VQ represent the same subspace. Hence the map

$$\Psi: \operatorname{St}(n,m) \longrightarrow \operatorname{Gr}(n,m), V \mapsto \operatorname{span}(V),$$

is an O(r)-quotient: two frames give the same subspace iff they differ by a right orthogonal factor.

Ellipsoidal Projection Metric. Following Definition 4, for $V, \bar{V} \in Gr(d, r)$,

$$\varrho_{\text{proj},\Sigma}(\mathcal{V},\bar{\mathcal{V}}) := \|\Sigma^{\frac{1}{2}}(\mathcal{P}_{\mathcal{V}} - \mathcal{P}_{\bar{\mathcal{V}}})\|_{\text{op}},\tag{59}$$

where $\mathcal{P}_{\mathcal{V}} := VV^{\top}$ for any V such that $\mathrm{span}(V) = \mathcal{V}$ (similarly $P_{\bar{\mathcal{V}}}$). Because $P_{\mathcal{V}}$ is unique for each subspace, $\varrho_{\mathrm{proj},\Sigma}$ is well defined (independent of the chosen V). The metric can be pulled back to $\mathrm{St}(n,m)$:

$$\varrho_{\operatorname{proj},\Sigma}(V,\bar{V}) := \varrho_{\operatorname{proj},\Sigma}(\operatorname{span}(V),\operatorname{span}(\bar{V})) = \|\Sigma^{\frac{1}{2}}(VV^{\top} - \bar{V}\bar{V}^{\top})\|_{\operatorname{op}}. \tag{60}$$

D.2 PRINCIPAL ANGLES BETWEEN SUB-SPACES

W study how metrics and angles between images \mathcal{V} and $\bar{\mathcal{V}}$ affect their spectral properties. We introduce principal angles and the cosine–sine (CS) decomposition—standard tools for analyzing subspaces (see, e.g., Chapter 6.4.3 in Golub & Van Loan (2013)).

Principle Angles and Cosine-Sine representation. Let U and \bar{U} be two $d \times d$ orthogonal matrix, and V and \bar{V} be the first r columns of U and \bar{U} , respectively. We are interested in studying the metrics and angles between r-dimensional sub-spaces $\mathcal{V} = \operatorname{span}(V)$ and $\bar{\mathcal{V}} = \operatorname{span}(\bar{V})$. Formally, denote

$$U, \ \bar{U} \in O(d), \qquad U = \begin{bmatrix} V & V_{\perp} \end{bmatrix}, \ \ \bar{U} = \begin{bmatrix} \bar{V} & \bar{V}_{\perp} \end{bmatrix},$$

where

$$V, \bar{V} \in \mathbb{R}^{d \times r}, \quad V^{\top}V = I_r, \quad \bar{V}^{\top}\bar{V} = I_r,$$

and

$$V_{\perp}, \bar{V}_{\perp} \in \mathbb{R}^{\,d \times (d-r)}, \quad V_{\perp}^{\top} V_{\perp} = \boldsymbol{I}_{\,d-r}, \quad \bar{V}_{\perp}^{\top} \bar{V}_{\perp} = \boldsymbol{I}_{\,d-r}.$$

Since $U, \bar{U} \in O(d)$, their product $U^{\top}\bar{U}$ is itself orthogonal. Writing

$$U^{\top} \bar{U} = \begin{pmatrix} V^{\top} \\ V_{\perp}^{\top} \end{pmatrix} \begin{bmatrix} \bar{V} & \bar{V}_{\perp} \end{bmatrix} = \begin{pmatrix} V^{\top} \bar{V} & V^{\top} \bar{V}_{\perp} \\ V_{\perp}^{\top} \bar{V} & V_{\perp}^{\top} \bar{V}_{\perp} \end{pmatrix},$$

define the four blocks

$$\underbrace{C}_{r \times r} = V^{\top} \bar{V}, \quad \underbrace{C_{\perp}}_{r \times (d-r)} = V^{\top} \bar{V}_{\perp}, \tag{61}$$

$$\underbrace{S}_{(d-r)\times r} = V_{\perp}^{\top} \bar{V}, \quad \underbrace{S_{\perp}}_{(d-r)\times (d-r)} = V_{\perp}^{\top} \bar{V}_{\perp}. \tag{62}$$

Thus

$$U^{\top}\bar{U} = \begin{pmatrix} C & C_{\perp} \\ S & S_{\perp} \end{pmatrix} \in O(d).$$

Now we introduce principal angles between $\mathcal{V} = \operatorname{span}(V)$ and $\bar{\mathcal{V}} = \operatorname{span}(\bar{V})$ by writing

$$C = V^{\top} \bar{V} = Q_1 \operatorname{diag}(\cos \theta_1, \cdots, \cos \theta_r) W_1^{\top}, \quad Q_1, W_1 \in O(r),$$
(63)

 $^{^1}$ By elementary linear algebra, the matrix definition of the orthogonal projector \mathcal{P} here coincides with the ℓ_2 -projection characterized in Lemma 14; thus the notation is consistent.

where

$$0 \le \theta_1 \le \theta_2 \le \dots \le \theta_r \le \pi/2$$

are called the principle angles between sub-spaces V and \bar{V} . Simultaneously, we have that the eigenvalues of S, C_{\perp} , S_{\perp} are (notation sepc means spectrum, the set of singular values)

$$spec(S) = \{ -\sin \theta_{1}, \cdots, -\sin \theta_{\min\{r, d-r\}}, \underbrace{0, \cdots, 0}_{\max\{d-2r, 0\}} \},$$

$$spec(C_{\perp}) = \{ \sin \theta_{1}, \cdots, \sin \theta_{\min\{r, d-r\}}, \underbrace{0, \cdots, 0}_{\max\{d-2r, 0\}} \}$$

$$spec(S_{\perp}) = \{ \cos \theta_{1}, \cdots, \cos \theta_{\min\{r, d-r\}}, \underbrace{1, \cdots, 1}_{\max\{d-2r, 0\}} \}.$$
(64)

The above representation in (63) and (64) are without loss of generality: if $r \le d-r$, then all the four spectrum contain all r principal angles; if r > d-r, then only first d-r principal angles $\{\theta_k\}_{k=1}^{d-r}$ can be smaller than $\pi/2$ and $\theta_k = 0$ for all $d-r+1 \le k \le r$.

The cosine–sine representation of the eigenvalues in (63) and (64) motivates our notation C and S when defining block matrices in (61) and (62). This representation is an immediate consequence of the classical CS decomposition for orthogonal matrices Paige & Wei (1994); Golub & Van Loan (2013), and we henceforth regard the resulting eigenvalue characterization as given.

Projection Metric via Principal Angles. For subspaces V and \bar{V} , recall that for orthogonal projectors

$$P_{\mathcal{V}} = VV^{\top}, \quad P_{\bar{\mathcal{V}}} = \bar{V}\bar{V}^{\top},$$

It is known that the projection metric defined in (59) and (60) are equal to $\sin \theta_r$, sine of the largest principal angle between the two sub-spaces. Formally, there is the fact (see, e.g., the last equation in Section 6.4.3 in Golub & Van Loan (2013))

$$\varrho_{\text{proj}} = \|P_{\mathcal{V}} - P_{\bar{\mathcal{V}}}\|_{\text{op}} = \max_{1 \le k \le r} \sin \theta_k = \sin \theta_r.$$
 (65)

Here θ_i is the *i*-th principal-angle between V and \overline{V} , and the spectral norm of the difference of two projectors equals the largest of these sines.

D.3 LOCAL CHARTS AND GLOBAL ATLAS OF THE GRASSMANNIAN

We present some facts about Grassmannian under the isotropic projection metric. These understanding will help us to establish Lemma 3 with the more challenging ellipsoidal projection metric.

In differential geometry, a *chart* is a single local coordinate map. An *atlas* is the whole collection of charts that covers the manifold. We introduce a useful atlas that consists of finite graph charts, which only rely on elementary linear algebra and avoid more advanced Lie algebra and exponential map techniques in Szarek (1997).

Choose a reference sub-space $\bar{\mathcal{V}}\in \mathrm{Gr}(d,r)$ and its parameterization matrix $\bar{\mathcal{V}}\in \mathrm{St}(d,r)$. Denote $X\in\mathbb{R}^{(d-r)\times r}$ to be mappings from r-dimensional subspace $\bar{\mathcal{V}}$ to (d-r)-dimensional subspace $\bar{\mathcal{V}}_{\perp}$. Every r-dimensional sub-space close to $\bar{\mathcal{V}}$ can be written as the graph

$$\mathcal{V}(X) := \operatorname{span}\left\{ \left[\bar{V} \bar{V}_{\perp} \right] \begin{pmatrix} I_r \\ X \end{pmatrix} \right\}, \qquad X \in \mathbb{R}^{(d-r) \times r}, \tag{66}$$

where is invariant to which parameterization matrix \bar{V} has been chosen. Given an appropriate local radius 1, define the local graph chart from $\{X \in \mathbb{R}^{(d-r)\times r} : \|X\|_{op} \leq 1\}$ to Gr(d,r) by

$$\phi_{\bar{\mathcal{V}}}: X \longmapsto \mathcal{V}(X) \in \operatorname{Gr}(d, r).$$
 (67)

Note that for the $(d-r) \times r$ zero matrix (denoted as 0), we have $\phi_{\bar{\nu}}(0) = \bar{\nu}$.

We will show that there is equivalence (up to absolute constant) between $\|X\|_{op}$ and $\varrho_{proj}(\mathcal{V}, \overline{\mathcal{V}})$. To be specific, we have the following lemma.

Lemma 21 (Sine-Tangent Representation in Graph Chart) Denote θ_r is the maximal principal angle between the sub-spaces V(X) and \bar{V} , defined in (63). For the graph chart (67), we have

$$\varrho_{\text{proj}}(\mathcal{V}(X), \bar{\mathcal{V}}) = \sin \theta_r, \quad ||X||_{\text{op}} = \tan \theta_r.$$

The above relationship immediately implies that

$$\varrho_{\text{proj}} = \|X\|_{\text{op}} / \sqrt{1 + \|X\|_{\text{op}}}.$$

We call X the graph parameterization of V(X) in this image.

Proof of Lemma 21: Given the fact $\varrho_{\text{proj}}(\mathcal{V}(X), \overline{\mathcal{V}}) = \sin \theta_r$ (which is already shown in (65)), where θ_r is the largest principal angle between the subspaces $\mathcal{V}(X)$ and the reference subspace $\overline{\mathcal{V}}$, we want to show $\|X\|_{\text{op}} = \tan \theta_r$.

Step 1: Setup and Simplification. The projection metric is invariant under orthogonal transformations of the ambient space \mathbb{R}^d . We can therefore choose a coordinate system that simplifies the calculations without loss of generality. We choose a basis such that the reference frame \bar{V} and its orthogonal complement \bar{V}_{\perp} are represented as:

$$\bar{V} = \begin{pmatrix} I_r \\ 0 \end{pmatrix} \in \operatorname{St}(d,r), \qquad \bar{V}_\perp = \begin{pmatrix} 0 \\ I_{d-r} \end{pmatrix} \in \operatorname{St}(d,d-r).$$

In this basis, the reference subspace is $\bar{\mathcal{V}} = \operatorname{span}(\bar{V})$. The parameterization matrix (orthonormal basis) V(X) for the subspace $\mathcal{V}(X)$ simplifies to (here $(I_r + X^\top X)^{-1/2}$ normalize V(X) to be an orthogonal matrix):

$$V(X) = [\bar{V} \ \bar{V}_{\perp}] \begin{pmatrix} I_r \\ X \end{pmatrix} (I_r + X^{\top} X)^{-1/2} = I_d \begin{pmatrix} I_r \\ X \end{pmatrix} (I_r + X^{\top} X)^{-1/2} = \begin{pmatrix} I_r \\ X \end{pmatrix} (I_r + X^{\top} X)^{-1/2}.$$

Step 2: Projection Metric and Principal Angles. A fundamental result in matrix analysis, our equation (63), states that the cosines of the principal angles, $\cos\theta_i$, between two subspaces spanned by orthonormal bases V and \bar{V} are the singular values of $V^\top \bar{V}$. In our case, the principal angles between V(X) and \bar{V} are determined by the singular values of $V(X)^\top \bar{V}$ —which are, equivalently, the singular values of $\bar{V}^\top V(X)$.

Step 3: Calculation of $\cos \theta_i$. Let's compute the matrix product $\bar{V}^\top V(X)$ using our simplified forms:

$$\bar{V}^{\top}V(X) = (I_r \quad 0) \left[\begin{pmatrix} I_r \\ X \end{pmatrix} (I_r + X^{\top}X)^{-1/2} \right]
= \left((I_r \quad 0) \begin{pmatrix} I_r \\ X \end{pmatrix} \right) (I_r + X^{\top}X)^{-1/2}
= I_r \cdot (I_r + X^{\top}X)^{-1/2}
= (I_r + X^{\top}X)^{-1/2}.$$

To find the singular values of this matrix, we use the Singular Value Decomposition (SVD) of X. Let $X = U\Sigma W^{\top}$, where $U \in \mathbb{R}^{(d-r)\times (d-r)}$ and $W \in \mathbb{R}^{r\times r}$ are orthogonal, and $\Sigma \in \mathbb{R}^{(d-r)\times r}$ is a rectangular diagonal matrix with the singular values $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$ on its diagonal. The spectral norm is $\|X\|_{\text{op}} = \lambda_1$.

Then, $X^{\top}X=(U\Sigma W^{\top})^{\top}(U\Sigma W^{\top})=W\Sigma^{\top}U^{\top}U\Sigma W^{\top}=W\Sigma_r^2W^{\top},$ where Σ_r^2 is the $r\times r$ diagonal matrix with entires λ_i^2 . So, the matrix $I_r+X^{\top}X=W(I_r+\Sigma_r^2)W^{\top}.$ Its inverse square root is: $(I_r+X^{\top}X)^{-1/2}=W(I_r+\Sigma_r^2)^{-1/2}W^{\top}.$

The singular values of $\bar{V}^{\top}V(X)$ are the diagonal entries of $(I_r + \Sigma_r^2)^{-1/2}$, which are: $s_i = \frac{1}{\sqrt{1+\lambda_i^2}}$.

These singular values are the values of $\cos \theta_i$. The largest principal angle, θ_r , corresponds to the smallest cosine value. This occurs when the singular value λ_i is largest, i.e., for $\lambda_1 = ||X||_{\text{op}}$. Thus,

$$\cos \theta_r = \frac{1}{\sqrt{1 + \|X\|_{\text{op}}^2}}.$$

Step 4: Deriving $\tan \theta_r$. Using the fundamental trigonometric identity $\sin^2 \theta + \cos^2 \theta = 1$ and the fact that principal angles lie in $[0.\pi/2]$, we have:

 $\tan \theta_r = ||X||_{\text{op}}.$

We have shown that for graph charts, there is the relationship $\varrho_{\text{proj}}(\mathcal{V}(X), \bar{\mathcal{V}}) = \sin \theta_r$ and $||X||_{\text{op}} = \tan \theta_r$. This suggests

$$\varrho_{\mathrm{proj}}(\mathcal{V}(X),\bar{\mathcal{V}}) = \frac{\|X\|_{\mathrm{op}}}{\sqrt{1+\|X\|_{\mathrm{op}}^2}}.$$

Global Finite Atlas of Graph Charts The following lemma proved in Pajor (1998) implies that an atlas consists of finite number of graph charts covers the entire Grassmannian manifold.

Lemma 22 (Finite atlas for Gr(d,r) and pointwise dimension consequence) For the Grassmannian Gr(d,r) we have that for all $\epsilon>0$, we have the coarse covering number bound $N(Gr(d,r),\varrho_{proj},\epsilon)\leq C^{\frac{r(d-r)}{\epsilon}}$, where C>0 is an absolute constant. This suggests that, a finite $O(e^{r(d-r)})$ number of graph charts are sufficient to cover the entire Gr(d,r) such that every sub-space $\mathcal{V}\in Gr(d,r)$ is contained in the image of a graph chart with its graph parameterization X satisfies $\|X\|_{op}\leq 1$. From this intuition, we have that the uniform prior $\mu=\mathrm{Unif}(Gr(d,r))$ satisfies that for every PSD Σ and $\epsilon>0$,

$$\log \frac{1}{\mu(B_{\varrho_{\mathrm{proj},\Sigma}}(\mathcal{V},\epsilon))} \leq C_1 r(d-r) + \sup_{X \in \bar{\mathcal{X}}} \log \frac{1}{\mathrm{Unif}(X' \in \bar{\mathcal{X}} : \varrho_{\mathrm{proj},\Sigma}(\mathcal{V}(X),\mathcal{V}(X')) \leq \epsilon)},$$

where $\mathcal{X}=\{X\in\mathbb{R}^{(d-r)r}:\|X\|_{\mathrm{op}}\leq 1\}$ and $\bar{\mathcal{X}}=\{X\in\mathbb{R}^{(d-r)r}:\|X\|_{\mathrm{op}}\leq 2\}$ (we make $\bar{\mathcal{X}}$ slightly larger for later technical derivation), and $C_1>0$ is an absolute constant.

Proof of Lemma 22: The preliminary, coarse covering number bound

$$N(Gr(d, r), \varrho_{proj}, \epsilon) \le C^{\frac{r(d-r)}{\epsilon}}$$

is stated as Proposition 6 in Pajor (1998), where C>0 is an absolute constant. By the definition of covering number (Definition 4.2.2 in Van Handel (2014)), we have that

$$N(Gr(d, r), \varrho_{proj}, \epsilon) \cdot Vol(\epsilon - \varrho_{proj} - ball) \ge Vol(Gr(d, r)),$$

then for the uniform prior $\nu = \text{Unif}(\text{Gr}(d,r))$, we have that for every $\bar{\mathcal{V}} \in \text{Gr}(d,r)$,

$$\log \frac{1}{\nu(B_{\varrho_{\mathrm{proj}}}(\bar{\mathcal{V}}, \epsilon))} = \log \frac{\mathrm{Vol}(\mathrm{Gr}(d, r))}{\mathrm{Vol}(\epsilon - \varrho_{\mathrm{proj}} - \mathrm{ball})} \leq r(d - r) \frac{\log C}{\epsilon}.$$

Taking $\epsilon = 1/\sqrt{2}$, we obtain:

$$\log \frac{1}{\nu(B_{\varrho_{\text{proj}}}(\bar{\mathcal{V}}, \sqrt{2}))} \le C_1 r(d-r),$$

where $C_1 > 0$ is an absolute constant. By Lemma 21, we have that inside the ball $B_{\varrho_{\text{proj}}}(\bar{\mathcal{V}}, 1/\sqrt{2})$, by choosing $\bar{\mathcal{V}}$ as the reference sub-space, the graph parameterization X of \mathcal{V} (see (66) for the definition) satisfies

$$||X||_{op} < 1.$$

Furthermore, $\mathcal{X} = \{X \in \mathbb{R}^{(d-r)r} : ||X||_{\text{op}} \leq 1\}$ satisfies (again by Lemma 21)

$$\mathcal{X} = B_{\varrho_{\mathrm{proj}}}(\bar{\mathcal{V}}, 1/\sqrt{2}) \subseteq \bar{\mathcal{X}} = B_{\varrho_{\mathrm{proj}}}(\bar{\mathcal{V}}, 2/\sqrt{5}).$$

Let

$$\mu_{\bar{\mathcal{V}}} = \mathrm{Unif}(B_{\mathrm{proj}}(\bar{\mathcal{V}}, 2/\sqrt{5})), \quad \mu(\mathcal{V}) = \int \nu(\bar{\mathcal{V}}) \mu_{\bar{\mathcal{V}}}(\mathcal{V}) d\bar{\mathcal{V}} = \mathrm{Unif}(\mathrm{Gr}(d, r)).$$

Then we have

$$\begin{split} \log \frac{1}{\mu(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V}, \epsilon))} &= \log \frac{1}{\int \nu(\bar{\mathcal{V}}) \mu_{\bar{\mathcal{V}}}(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V}, \epsilon)) d\bar{\mathcal{V}}} \\ &= \log \frac{1}{\int \nu(\bar{\mathcal{V}}) \mu_{\bar{\mathcal{V}}}(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V}, \epsilon) \cap B_{\text{proj}}(\bar{\mathcal{V}}, 2/\sqrt{5})) d\bar{\mathcal{V}}} \\ &\leq \log \frac{1}{\nu(B_{\varrho_{\text{proj}}}(\mathcal{V}, 1/\sqrt{2})) \min_{\bar{\mathcal{V}} \in B_{\varrho_{\text{proj}}}(\mathcal{V}, 1/\sqrt{2})} \mu_{\bar{\mathcal{V}}}(X' \in \bar{\mathcal{X}} : \varrho_{\text{proj},\Sigma}(\mathcal{V}(X), \mathcal{V}(X')) \leq \epsilon)} \\ &\leq C_1 r(d-r) + \sup_{X \in \mathcal{X}} \log \frac{1}{\operatorname{Unif}(X' \in \bar{\mathcal{X}} : \varrho_{\text{proj},\Sigma}(\mathcal{V}(X), \mathcal{V}(X')) \leq \epsilon)}, \end{split}$$

where the first inequality is by restricting $\bar{\mathcal{V}}$ to $B_{\varrho_{\text{proj}}}(\mathcal{V},1/\sqrt{2})$. Note that we use different radius here than in $\mu_{\bar{\mathcal{V}}}$ to enusre that the set $\bar{\mathcal{X}}$ for X', which is inside the uniform distribution in the final bound, to be larger than the domain \mathcal{X} for X to take sup. This will help later technical derivation. \square

D.4 LIPCHITZ OF GRAPH CHART

We have the following Lipchitz property of graph chart.

Lemma 23 (Lipchitz of Graph Chart) Let $X, X' \in \mathbb{R}^{(d-r)\times r}$ be matrices that satisfies $||X|| \leq 1$ and $||X'|| \leq 2$. Define the function $A : \mathbb{R}^{(d-r)\times r} \to \mathbb{R}^{d\times r}$ and $M : \mathbb{R}^{(d-r)\times r} \to \mathbb{R}^{r\times r}$ by

$$A(X) = \begin{pmatrix} I_r \\ X \end{pmatrix}, \quad M(X) = (I_r + X^\top X)^{-1/2}.$$

Then the ellipsoidal projection metric is Lipschitz to ellipsoidal metrics as follows: for every rank- $PSD \Sigma_1 \in \mathbb{R}^{d \times d}$,

$$\begin{split} &\varrho_{\text{proj}}(\mathcal{V}(X), \mathcal{V}(X')) \\ = &\| \Sigma^{\frac{1}{2}}(A(X)M(X)A(X)^{\top} - A(X')M(X')A(X')^{\top}) \|_{\text{op}} \\ \leq &\sqrt{5} \| \left((0 \quad I_{d-r}) \Sigma \begin{pmatrix} 0 \\ I_{d-r} \end{pmatrix} \right)^{\frac{1}{2}} (X - X') \|_{\text{op}} + \sqrt{5} \| \Sigma^{\frac{1}{2}}A(X)M(X)X^{\top}(X' - X) \|_{\text{op}} \\ &+ (1 + \sqrt{5}) \| \Sigma^{\frac{1}{2}}A(X)M(X)(X' - X)^{\top} \|_{\text{op}}. \end{split}$$

We continue to present the following lemma, which shows that the A(X) and M(X) factors in Lemma 23 only reduces the effective dimensions of the ellipsoidal map, and does not increase the eigenvalues (up to absolute constants).

Proof of Lemma 23: We use the triangle inequality:

$$\begin{split} &A(X)M(X)A(X) - A(X')M(X')A(X') \\ = &(A(X) - A(X'))M(X')A(X') + A(X)(M(X) - M(X'))A(X') + A(X)M(X)(A(X) - A(X')). \end{split}$$
 Therefore,

$$\|\Sigma^{\frac{1}{2}}(A(X)M(X)A(X)^{\top} - A(X')M(X')A(X')^{\top})\|_{\text{op}}$$

$$\leq \|\Sigma^{\frac{1}{2}} \begin{pmatrix} 0 \\ X - X' \end{pmatrix} M(X')A(X')^{\top}\|_{\text{op}} + \|\Sigma^{\frac{1}{2}}A(X)(M(X) - M(X'))A(X')^{\top}\|_{\text{op}}$$

$$+ \|\Sigma^{\frac{1}{2}}A(X)M(X) \begin{pmatrix} 0 & X^{\top} - {X'}^{\top} \end{pmatrix}\|_{\text{op}}.$$
(68)

We will bound each term.

First, by singular decomposition of X' and $\|X'\|_{op} \le 2$, we know $\|A(X')\|_{op} \le \sqrt{5}$, $\|M(X')\|_{op} \le 1$, and

$$\begin{split} &\| \Sigma^{\frac{1}{2}} \begin{pmatrix} 0 \\ X - X' \end{pmatrix} M(X') A(X')^{\top} \|_{\text{op}} \leq \sqrt{2} \| \Sigma^{\frac{1}{2}} \begin{pmatrix} 0 \\ X - X' \end{pmatrix} \|_{\text{op}} \\ &= \sqrt{5} \| \Sigma^{\frac{1}{2}} \begin{pmatrix} 0 \\ I_{d-r} \end{pmatrix} (X - X') \|_{\text{op}} = \sqrt{2} \| \begin{pmatrix} (0 & I_{d-r}) \Sigma \begin{pmatrix} 0 \\ I_{d-r} \end{pmatrix} \end{pmatrix}^{\frac{1}{2}} (X - X') \|_{\text{op}}. \end{split}$$

Second, because

$$M(X) - M(X')$$

$$= (I_r + X^{\top}X)^{-1} \left((I_r + {X'}^{\top}X') - (I_r + X^{\top}X) \right) (I_r + {X'}^{\top}X')^{-1}$$

$$= (I_r + X^{\top}X)^{-1} \left({X'}^{\top}X' - X^{\top}X \right) (I_r + {X'}^{\top}X')^{-1}$$

$$= (I_r + X^{\top}X)^{-1} \left({X^{\top}(X' - X) + ({X'}^{\top} - X^{\top})X'} \right) (I_r + {X'}^{\top}X')^{-1}$$

$$= M(X)X^{\top}(X' - X)M(X') + M(X)({X'}^{\top} - X^{\top})X'M(X'),$$

we have

$$\begin{split} & \| \Sigma^{\frac{1}{2}} A(X) (M(X) - M(X')) A(X')^{\top} \|_{\text{op}} \\ \leq & \sqrt{5} \| \Sigma^{\frac{1}{2}} A(X) M(X) X^{\top} (X' - X) \|_{\text{op}} + \sqrt{2} \| \Sigma^{\frac{1}{5}} A(X) M(X) (X' - X)^{\top} \|_{\text{op}}. \end{split}$$

Third, we have

$$\|\Sigma^{\frac{1}{2}}A(X)M(X)\left(0 \quad X^{\top} - {X'}^{\top}\right)\|_{\text{op}} = \|\Sigma^{\frac{1}{2}}A(X)M(X)(X - X')^{\top}\|_{\text{op}},$$

Substituting these bounds back into (68), we prove the desired result.

Lemma 24 (Spectral domination under contractions) Let $\Sigma \succeq 0$ be a $d \times d$ PSD matrix with ordered eigenvalues $\lambda_1(\Sigma) \geq \cdots \geq \lambda_d(\Sigma)$. Let $A \in \mathbb{R}^{d \times r}$ and write $s := ||A||_{\text{op}}$. Denote by $\mu_1 \geq \cdots \geq \mu_r$ the eigenvalues of $A^{\top} \Sigma A$. Then, for every $k = 1, \ldots, r$,

 $\mu_k \leq s^2 \lambda_k(\Sigma).$

Proof of Lemma 24: By the Courant–Fischer–Weyl max-min characterization (see, e.g., Wikipedia contributors (b)), we have

$$\lambda_{k}(A^{\top}\Sigma A) = \min_{\substack{S \subset \mathbb{R}^{d} \\ \dim S = d - k + 1}} \sup\{\|A^{\top}\Sigma^{\frac{1}{2}}x\|_{2}^{2} : x \in S, \|x\|_{2} = 1\}$$

$$\leq s^{2} \cdot \min_{\substack{S \subset \mathbb{R}^{d} \\ \dim S = d - k + 1}} \sup\{\|\Sigma^{1/2}x\|_{2} : x \in S, \|x\|_{2} = 1\}$$

$$= s^{2}\lambda_{k}(\Sigma).$$

D.5 PROOF OF THE MAIN RESULT

From Lemma 22, to cover $\operatorname{Gr}(d,r)$ it suffices to cover the unit ball of $(d-r)\times r$ matrices under the ellipsoidal spectral metric. We are now ready to prove Lemma 3, our main result for ellipsoidal Grassmannian covering.

Proof of Lemma 3: Define $\mathcal{X}=\left\{X\in\mathbb{R}^{(d-r)\times r}:\|X\|_{\mathrm{op}}\leq 1\right\}$ and $\bar{\mathcal{X}}=\left\{X\in\mathbb{R}^{(d-r)\times r}:\|X\|_{\mathrm{op}}\leq 2\right\}$. By Lemma 22 (Global Property about Finite Atlas), we have that for all $\epsilon>0$,

$$\log \frac{1}{\mu(B_{\varrho_{\text{proj},\Sigma}}(\mathcal{V},\epsilon))} \le C_1 r(d-r) + \sup_{X \in \mathcal{X}} \log \frac{1}{\text{Unif}(X' \in \bar{\mathcal{X}} : \varrho_{\text{proj},\Sigma}(\mathcal{V}(X),\mathcal{V}(X')) \le \epsilon)}, \quad (69)$$

where $C_1 > 0$ is an absolute constant.

Given a PSD matrix $H \in \mathbb{R}^{p \times p}$ and an eigenvalue threshold α , assume its eigendecomposition is $H = U \operatorname{diag}(\beta_1, \cdots, \beta_p)U^{\top}$, define the thresholding function T_{α} by

$$T_{\alpha}(H) = U \operatorname{diag}(\max\{\beta_1, \alpha\}, \cdots, \max\{\beta_p, \alpha\})U^{\top}.$$

Clearly this function only increases the metric. Consider the $d \times (d-r)$ matrix $A(X) = \begin{pmatrix} I_r \\ X \end{pmatrix}$

and the $r \times e$ matrix $M(X) = (I_r + X^{\top}X)^{-1}$, we are interested in the following four ellipsoidal metric, where

$$\begin{split} \varrho_1^2(X,X') = & 6(X'-X)^\top T_{\epsilon^2/42} \left((0 \quad I_{d-r}) \, \Sigma \begin{pmatrix} 0 \\ I_{d-r} \end{pmatrix} \right) (X'-X), \\ \varrho_2^2(X,X') = & 6(X'-X)^\top T_{\epsilon^2/21} \left(X M(X) A(X)^\top \Sigma A(X) M(X) X^\top \right) (X'-X), \\ \varrho_3^3(X,X') = & 12(X'-X) T_{\epsilon^2/21} \left(M(X) A(X)^\top \Sigma A(X) M(X) \right) (X-X')^\top. \end{split}$$

By Lemma 23 (Lipchitz of Graph Chart), ellipsoidal projection metric is bounded by $\varrho_1 + \varrho_2 + \varrho_3$, so for any $X \in \mathcal{X}$,

$$\begin{split} \log \frac{1}{\operatorname{Unif}(X' \in \mathcal{X} : \varrho_{\operatorname{proj},\Sigma}(\mathcal{V}(X), \mathcal{V}(X')) \leq \epsilon)} \\ \leq \log \frac{1}{\operatorname{Unif}(X' \in \bar{\mathcal{X}} : \varrho_{1}(X, X') + \varrho_{2}(X, X') + \varrho_{3}(X, X') \leq \epsilon)} \\ \leq \log \frac{1}{\operatorname{Unif}(X' \in \bar{\mathcal{X}} : \varrho_{1}(X, X') \leq \frac{\epsilon}{\sqrt{7}}) \cdot \operatorname{Unif}(X' \in \mathcal{X} : \varrho_{2}(X, X') \leq \frac{\sqrt{2}\epsilon}{\sqrt{7}})) \cdot \operatorname{Unif}(X' \in \bar{\mathcal{X}} : \varrho_{3}(X, X') \leq \frac{\sqrt{4}\epsilon}{\sqrt{7}})))} \\ = \log \frac{1}{\operatorname{Unif}(B_{\varrho_{1}}(X, \epsilon/\sqrt{7}))} + \log \frac{1}{\operatorname{Unif}(B_{\varrho_{2}}(X, \sqrt{2}\epsilon/\sqrt{7}))} + \log \frac{1}{\operatorname{Unif}(B_{\varrho_{3}}(X, \sqrt{4}\epsilon/\sqrt{7}))}, \quad (70) \end{split}$$

where the last equality is due to the fact that all the three balls $B_{\varrho_1}(X,\epsilon/\sqrt{7})$, $B_{\varrho_2}(X,\sqrt{2}\epsilon/\sqrt{7})$, $B_{\varrho_3}(X,\sqrt{4}\epsilon/\sqrt{7})$ are contained in $\bar{\mathcal{X}}$, as we have applied the thresholding function to ensure this inclusion. For example, for the first ball, from

$$X' - X = \left(\underbrace{\sqrt{6}T_{\epsilon^2/42}\left(\begin{pmatrix} 0 & I_{d-r}\end{pmatrix}\Sigma\begin{pmatrix} 0 \\ I_{d-r}\end{pmatrix}\right)}_{\text{denoted as } H_1}\right)^{-1/2} \sqrt{6}T_{\epsilon^2/42}\left(\begin{pmatrix} 0 & I_{d-r}\end{pmatrix}\Sigma\begin{pmatrix} 0 \\ I_{d-r}\end{pmatrix}\right)^{\frac{1}{2}} (X' - X)$$

we have (combining with $\lambda_{\min}(H_1) \geq \epsilon^2/7$)

$$||X' - X||_{\text{op}} \le \lambda_{\min}(H_1)^{-1/2} \cdot \epsilon / \sqrt{7} \le 1,$$

which resulting in $||X'||_{op} \le ||X' - X||_{op} + ||X||_{op} \le 2$.

A classical volume-ratio argument (e.g., display (1) in Pajor (1998); see also Milman & Schechtman (1986); Pisier (1999)) states the following result on the covering number of balls in general normed space: for a p-dimensional normed space $\mathcal Y$ equipped with the metric associated to its norm $\|\cdot\|$, we denote by B(0,R) the ball in $\mathcal Y$ centered at origin with radius R, and by $N(B(0,R),\|\cdot\|,\epsilon)$ the covering number of this ball. Then for every $\varepsilon\in(0,R]$,

$$\left(\frac{R}{\varepsilon}\right)^p \; \leq \; \mathrm{N}\!\big(B(0,R),\, \lVert \cdot \rVert,\, \varepsilon\big) \; \leq \; \left(\frac{3R}{\varepsilon}\right)^p.$$

Note that this result is for general normed space, not only for the ℓ_2 norm in Euclidean space. By the definition of the covering number (Definition 4.2.2 in Van Handel (2014)), we have that for every $X \in \mathcal{X}$ and every $\epsilon \in (0,1]$,

$$\frac{\operatorname{Vol}(\bar{\mathcal{X}})}{\operatorname{Vol}(B_{\operatorname{op}}(X,\epsilon))} \ \leq \ \operatorname{N}\!\!\left(B_{\operatorname{op}}(X,\epsilon), \, \lVert \cdot \rVert_{\operatorname{op}}, \, \varepsilon\right) \ \leq \ \left(\frac{6}{\varepsilon}\right)^p.$$

Applying the standard change of variable argument $Y=H_1^{\frac{1}{2}}X$ and Jacobian determinant calculation, we have

$$\operatorname{Vol}(B_{\varrho_1}(X,\varepsilon/\sqrt{7})) = \operatorname{Vol}(B_{\operatorname{op}}(X,\varepsilon/\sqrt{7})) (\det H_1)^{-1/2} = \operatorname{Vol}(B_{\operatorname{op}}(X,\varepsilon/\sqrt{7})) \prod_{k=1}^p \lambda_k(H_1)^{-1/2},$$

which gives (again using the fact $B_{\rho_1}(X, \epsilon/\sqrt{7}) \subseteq \bar{\mathcal{X}})$

$$\log \frac{1}{\mathrm{Unif}(B_{\varrho_1}(X,\epsilon/\sqrt{7})} \le r \sum_{k=1}^{d-r} \log \frac{42\lambda_k(H)}{\epsilon^2} \le \frac{1}{2}r \sum_{k=1}^{d-r} \log \frac{C_2 \max\{\lambda_k(\Sigma), \epsilon^2/42\}}{\epsilon^2},$$

where the last inequality uses Lemma 24, and $C_2 > 0$ is an absolute constant. Similarly, we have

$$\log \frac{1}{\mathrm{Unif}(B_{\varrho_2}(X,\sqrt{2}\epsilon/\sqrt{7})} \leq \frac{1}{2}r\sum_{k=1}^{d-r}\log \frac{C_3\max\{\lambda_k(\Sigma),\epsilon^2/42\}}{\epsilon^2},$$

where $C_3 > 0$ is an absolute constant, and

$$\log \frac{1}{\mathrm{Unif}(B_{\varrho_3}(X,\sqrt{4}\epsilon/\sqrt{7})} \leq \frac{1}{2}(d-r)\sum_{k=1}^r \log \frac{C_4 \max\{4\lambda_k(\Sigma),\epsilon^2/42\}}{\epsilon^2},$$

where $C_4 > 0$ is an absolute constant

Combining the above three inequalities with (69) (70), we have that for all $\epsilon > 0$, we have

$$\log \frac{1}{\mu(B_{\varrho_{\text{proj}},\Sigma}(\mathcal{V},\epsilon))} \le \frac{1}{2}(d-r)\sum_{k=1}^{r}\log \frac{C\max\{\lambda_{k},\epsilon^{2}/42\}}{\epsilon^{2}} + r\sum_{k=1}^{d-r}\log \frac{C\max\{\lambda_{k},\epsilon^{2}/42\}}{\epsilon^{2}},\tag{71}$$

where C > 0 is an absolute constant.

E PROOFS FOR GENERALIZATION BOUNDS AND COMPARISON (SECTION 4)

E.1 Proof of Theorem 4 in Section 4.1

The proof consists of two steps: 1. Obtaining the Integral Bound on Generalization Gap; and 2. Obtaining the Expression of Riemannian Dimension.

Step 1: Obtaining the Integral Bound on Generalization Gap. As presented in (6), we construct the metric tensor

$$G_{\mathrm{NP}}(W) := \mathrm{blockdiag}\left(\cdots, LM_{l \to L}^2(W, \epsilon) \cdot F_{l-1}(W, X) F_{l-1}^{\top}(W, X) \otimes I_{d_l}, \cdots\right).$$

By Lipchitz property of the loss function we have

$$\varrho_{n,\ell}(W',W) = \sqrt{\mathbb{P}_n(\ell(W';(x,y)) - \ell(W;(x,y)))^2}$$

$$\leq \beta \sqrt{\mathbb{P}_n \|f(W',x) - f(W,x)\|_2^2} = \beta \varrho_n(W',W)$$

By Lemma 1 we have the metric dominating relationship: for every $W \in \mathcal{W}$,

$$\sqrt{n}\varrho_n(W',W) \le \varrho_{G_{NP}(W)}(W',W), \quad \forall W' \in \mathcal{W}.$$

Combining the above two inequalities we have

$$\varrho_{n,\ell}(W',W) \le \frac{\beta}{\sqrt{n}} \varrho_{G_{NP}(W)}(W',W), \quad \forall W' \in \mathcal{W}.$$

By the metric domination lemma (Lemma 13), we have the pointwise dimension bound: for every $w \in \mathcal{W}$,

$$\log \frac{1}{\pi(B_{\rho_{n,\ell}}(W,\sqrt{n}\epsilon/\beta))} \leq \log \frac{1}{\pi(B_{G_{\mathrm{NP}}(W)}(W,\sqrt{n}\epsilon/\beta))}.$$

By definition of Riemannian dimension, if the manifold (W, G) admits Riemannian dimension d_R , then we have that for every $W \in W$,

$$\log \frac{1}{\pi(B_{\varrho_{n,\ell}}(W,\sqrt{n}\epsilon/\beta))} \le \log \frac{1}{\pi(B_{G_{NP}(W)}(W,\sqrt{n}\epsilon/\beta))} \le d_R(W,\epsilon/\beta). \tag{72}$$

By Theorem 2, we have that there exists an absolute constant C_1 such that with probability at least $1 - \delta$, uniformly over all $W \in \mathcal{W}$,

$$(\mathbb{P} - \mathbb{P}_{n})\ell(f(W, x), y) \leq C_{1} \left(\frac{1}{\sqrt{n}} \int_{0}^{\infty} \sqrt{\log\left(\frac{1}{\pi(B_{\varrho_{n,\ell}}(W, \epsilon))}\right)} d\epsilon + \sqrt{\frac{\log\frac{\log(2n)}{\delta}}{n}} \right)$$

$$\leq C_{1} \left(\frac{1}{\sqrt{n}} \int_{0}^{\infty} \sqrt{d_{R}(W, \epsilon/\beta)} d\epsilon + \sqrt{\frac{\log\frac{\log(2n)}{\delta}}{n}} \right)$$

$$= C_{1} \left(\frac{\beta}{\sqrt{n}} \int_{0}^{\infty} \sqrt{d_{R}(W, \epsilon)} d\epsilon + \sqrt{\frac{\log\frac{\log(2n)}{\delta}}{n}} \right). \tag{73}$$

where C_1 is an absolute constant; the first inequality uses Theorem 2; and the second inequality uses (72). This finishes the first part of Theorem 4 (integral upper bound).

Step 2: Obtaining the Expression of Riemannian Dimension. It remains to express the Riemannian dimension d_R by Theorem 3 and prove the second part of Theorem 4. By Theorem 3, we have that the manifold (W, G) admits the Riemannian dimension

$$d_{R}(W,\epsilon) = \sum_{l=1}^{L} \left((d_{l} + d_{l-1}) \cdot d_{\text{eff}}(LM_{l \to L}^{2}(W,\epsilon) \cdot F_{l-1}(W,X) F_{l-1}(W,X)^{\top}, C_{2} \max\{\|W\|_{\mathbf{F}}, 1/n\}, \epsilon) + \log(d_{l-1}n) \right),$$
(74)

where $R = \sup_{\mathcal{W}} \|W\|_{\mathbf{F}}$, C_2 is an absolute constant, and the effective dimension (defined via (8)) is

$$d_{\text{eff}}(LM_{l\to L}^{2}(W,\epsilon) \cdot F_{l-1}(W,X)F_{l-1}(W,X)^{\top}, C_{2} \max\{\|W\|_{\mathbf{F}}, 1/n\}, \epsilon)$$

$$= \frac{1}{2} \sum_{k=1}^{r_{\text{eff}}[W,l]} \log \frac{8C_{2}^{2} \max\{\|W\|_{\mathbf{F}}^{2}, R^{2}/4^{n}\}LM_{l\to L}^{2}(W,\epsilon)\lambda_{k}(F_{l-1}F_{l-1})^{\top}}{n\epsilon^{2}},$$
(75)

where F_{l-1} is the abbreviation of $F_{l-1}(W,X)$ and $r_{\mathrm{eff}}[W,l]$ is the abbreviation of $r_{\mathrm{eff}}(LM^2_{l\to L}(W,\epsilon)\cdot F_{l-1}(W,X)F_{l-1}(W,X)^\top, C_2\max\{\|W\|_{\mathbf{F}},R/2^n\},\epsilon).$

Combining the identities (74) and (75), we have the pointwise dimension bound

$$d_{R}(W,\epsilon) = \sum_{l=1}^{L} \left((d_{l} + d_{l-1}) \sum_{k=1}^{r_{\text{eff}}[W,l]} \log \frac{8C_{2}^{2} \lambda_{k} (F_{t-1} F_{t-1}^{\top}) \cdot \max\{\|W\|_{\mathbf{F}}^{2}, R^{2} / 4^{n}\} L M_{l \to L}(W,\epsilon)}{n\epsilon^{2}} + \log(d_{l-1}n) \right),$$
(76)

where F_{l-1} is the abbreviation of $F_{l-1}(W,X)$; $r_{\text{eff}}[W,l]$ is the abbreviation of $r_{\text{eff}}(LM_{l\to L}^2(W,\epsilon) \cdot F_{l-1}(W,X)F_{l-1}(W,X)^\top, C_2 \max\{\|W\|_{\mathbf{F}}, 1/n\}, \epsilon)$; and C_2 is an absolutant.

This finishes the second part of Theorem 4 (expression of Riemannian dimension).

Combining the integral upper bound (73) and the Riemannain dimension expression (76) concludes the proof of Theorem 4.

E.2 PROOF OF IMPLICIT BIAS IN SECTION 4.2

For the regularized ERM $\hat{f} = \operatorname{argmin}_f \{ \mathbb{P}_n \ell(f; z) + C \sqrt{\frac{d(f) + \log(1/\delta)}{n}} \}$, with probability at least $1 - \delta$

$$\mathbb{P}\ell(\hat{f};z) \leq \inf_{f \in \mathcal{F}} \left\{ \mathbb{P}_n \ell(f;z) + C\sqrt{\frac{d(f) + \log(1/\delta)}{n}} \right\}
\leq \inf_{f \in \mathcal{F}} \left\{ \mathbb{P}\ell(f;z) + 2C\sqrt{\frac{d(f) + \log(1/\delta)}{n}} \right\},$$
(77)

where the first inequality uses the bound of the form (3) and the second inequality uses the following result of (78). Define $\ell'(f;z) := 1 - \ell(f;z) \in [0,1]$. Applying the bound in (3) with ℓ' , we get a generalization error bound

$$(\mathbb{P}_n - \mathbb{P})\ell(f, z) \le C\sqrt{\frac{d(f) + \log(1/\delta)}{n}},\tag{78}$$

where the complexity measure remains unchanged, which can be verified by the pointwise dimension. The bound in (77) then implies

$$\mathbb{P}\ell(\hat{f};z) - \mathbb{P}\ell(f^*;z) \le \inf_{f \in \mathcal{F}} \left\{ \mathbb{P}\ell(f;z) + 2C\sqrt{\frac{d(f) + \log(1/\delta)}{n}} \right\} - \mathbb{P}\ell(f^*;z),$$

which, in turns, means that

$$\mathbb{P}\ell(\hat{f};z) - \mathbb{P}\ell(f^*;z) \le 2C\sqrt{\frac{d(f^*) + \log(1/\delta)}{n}}.$$

E.3 Proof of the Improvement over Norm Bounds in Section 4.3

We now provide norm-constrained bound from Theorem 4 without any expression $r_{\rm eff}$ and $d_{\rm eff}$ in the bound. Invoking the elementary bound $\log x \leq \log(1+x) \leq x$, the effective dimension factor in Theorem 4 can be relaxed to the dimension-independent bound

$$\sum_{k=1}^{\infty} \log \left(\frac{\lambda_{k} \left(F_{l-1} F_{l-1}^{\top} \right) \|W\|_{\mathbf{F}}^{2} L M_{l \to L}^{2}(W, \varepsilon)}{n \varepsilon^{2}} \right) \leq \frac{\sum_{k=1}^{\infty} \lambda_{k} \left(F_{l-1} F_{l-1}^{\top} \right) \|W\|_{\mathbf{F}}^{2} L M_{l \to L}^{2}(W, \varepsilon)}{n \varepsilon^{2}}$$

$$\leq \frac{\|F_{l-1}(W, X)\|_{\mathbf{F}}^{2} \|W\|_{\mathbf{F}}^{2} L M_{l \to L}^{2}(W, \varepsilon)}{n \varepsilon^{2}},$$

and one arrives at the following rank–free consequence.

Corollary 1 (Norm-constrained bound) Theorem 4 is never worse than: uniformly over all $W \in \mathcal{W}$, the generalization gap $(\mathbb{P} - \mathbb{P}_n) \ell(f(W, x), y)$ is bounded by

$$O\left(\frac{\beta\sqrt{\sum_{l=1}^{L}(d_{l}+d_{l-1})L\|F_{l-1}(W,X)\|_{\mathbf{F}}^{2}\|W\|_{\mathbf{F}}^{2}\sup_{\epsilon>0}M_{l\to L}^{2}(W,\epsilon)}{n}+\sqrt{\frac{\beta^{2}\sum_{l=1}^{L}\log(d_{l-1}n)+\log\frac{\log(2n)}{\delta}}{n}}\right)$$
(79)

Furthermore, (79) implies the spectrally normalized bound: uniformly over W, the generalization gap $(\mathbb{P} - \mathbb{P}_n) \ell(f(W, x), y)$ is bounded by

$$O\left(\frac{\beta \|X\|_{\mathbf{F}} \|W\|_{\mathbf{F}} \cdot \sqrt{\sum_{l=1}^{L} L(d_{l} + d_{l-1}) \prod_{i \neq l} \|W_{i}\|_{\text{op}}^{2}}}{n} + \sqrt{\frac{\beta^{2} \sum_{l=1}^{L} \log(d_{l-1}n) + L \log \frac{n \log \max\{R, 2\}}{\delta}}{n}}\right)$$
(80)

Here in both (79) and (80), O hides multiplicative absolute constants and two ignorable high-order terms: $\frac{\beta\sqrt{\sum_{l=1}^{L}(d_{l}+d_{l-1})d_{l-1}}}{n^{5.5}} \text{ and } \frac{\beta\sqrt{\sum_{l=1}^{L}(d_{l}+d_{l-1})L\|F_{l-1}\|_{\mathsf{F}}^{2}R^{2}\sup_{\epsilon>0}M_{l\to L^{2}(W,\epsilon)}}}{n^{2n}}; \text{ and in (80),}$ O additionally hides an ignorable high-order term $\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2}\sum_{l=1}^{L}(d_{l}+d_{l-1})(R/\sqrt{L-1})^{L-1}}}{n\max\{R,2\}^{n}}.$

Remark: in (79) and (80), the feature matrices $F_{l-1}(W;X)$ and X contain n features vectors so their Frobenius norms scales with \sqrt{n} , making the order of both bounds to be $n^{-1/2}$.

Proof of Corollary 1. Riemannain Dimension can be expressed in the following equivalent form

$$\int_0^\infty \sqrt{d_R(W,\epsilon)} d\epsilon = \inf_{\alpha > 0} \left(\int_0^\alpha \sqrt{d_R(W,\epsilon)} d\epsilon + \int_\alpha^1 \sqrt{d_R(W,\epsilon)} d\epsilon \right).$$

We organize the proof with four steps.

Step 1: Bounding the Dominating Integral. As we will take α to be very small so that the $\int_0^\alpha \sqrt{d_R(W,\epsilon)}$ will be not exceed the order of $\int_\alpha^\infty \sqrt{d_R(W,\epsilon)}$, we firstly prove $\int_\alpha^1 \sqrt{d_R(W,\epsilon)} d\epsilon$. By the basic inequality $\log x \leq \log(1+x) \leq x$ for x>0, we have

$$\sum_{k=1}^{r_{\text{eff}}[W,l]} \log \left(\frac{8C_{2}^{2}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2}, R^{2}/4^{n}\}LM_{l\to L}^{2}(W,\epsilon)}{n\epsilon^{2}} \right) \\
\leq \sum_{k=1}^{r_{\text{eff}}[W,l]} \frac{8C_{2}^{2}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2}, R^{2}/4^{n}\}LM_{l\to L}^{2}(W,\epsilon)}{n\epsilon^{2}} \\
\leq \sum_{k=1}^{d_{l-1}} \frac{32C_{2}^{2}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2}, R^{2}/4^{n}\}LM_{l\to L}^{2}(W,\epsilon)}{n\epsilon^{2}} \\
= \frac{8C_{2}^{2}||F_{l-1}||_{F}^{2}\{||W||_{F}^{2}, R^{2}/4^{n}\}LM_{l\to L}^{2}(W,\epsilon)}{n\epsilon^{2}}, \tag{81}$$

where F_{l-1} is the abbreviation of $F_{l-1}(W,X)$; $r_{\text{eff}}[W,l]$ is the abbreviation of $r_{\text{eff}}(LM_{l\to L}^2(W,\epsilon) \cdot F_{l-1}(W,X)F_{l-1}(W,X)^\top$, $C_2 \max\{\|W\|_{\textbf{F}}, 1/n\}, \epsilon\}$; and C_2 is a positive absolute constant. Here the second inequality uses the definition that $r_{\text{eff}}[W,l]$ as the effective rank of a $d_{l-1} \times d_{l-1}$ matrix, is no larger than the matrix width d_{l-1} ; the first equality is because

$$\sum_{k=1}^{d_{l-1}} \lambda_k(F_{l-1}F_{l-1}^{\top}) = \text{Tr}(F_{l-1}F_{l-1}^{\top}) = ||F_{l-1}||_{\mathbf{F}}^2, \tag{82}$$

a well-known property of the Frobenius norm (the squared Frobenius norm $||F_{l-1}||_F^2$ equals trace of $F_{l-1}F_{l-1}^{\top}$). By (81) and Theorem 4 we have the Riemannian Dimension upper bound

$$d_R(W,\epsilon) \le 8C_2^2 \sum_{l=1}^{L} (d_l + 2d_{l-1}) \frac{\|F_{l-1}\|_{\mathbf{F}}^2 \max\{\|W\|_{\mathbf{F}}^2, R^2/4^n\} L M_{l \to L}^2(W,\epsilon)}{n\epsilon^2} + \sum_{l=1}^{L} \log(d_{l-1}n),$$
(83)

where C_2 is a positive absolute constant.

Taking (83) to the integral $\int_{\alpha}^{1} \sqrt{d_R(W,\epsilon)} d\epsilon$, we have

$$\begin{split} & \int_{\alpha}^{1} \sqrt{d_{R}(W,\epsilon)} d\epsilon \\ \leq & 2\sqrt{2}C_{2} \int_{\alpha}^{1} \sqrt{\sum_{l=1}^{L} (d_{l} + 2d_{l-1}) \frac{\|F_{l-1}\|_{\mathbf{F}}^{2} \max\{\|W\|_{\mathbf{F}}^{2}, R^{2}/4^{n}\} L M_{l \to L}^{2}(W,\epsilon)}{n\epsilon^{2}}} d\epsilon + (1-\alpha) \sqrt{\sum_{l=1}^{L} \log(d_{l-1}n)} \\ \leq & C_{3} \sqrt{\frac{\sum_{l=1}^{L} (d_{l} + d_{l-1}) L \|F_{l-1}\|_{\mathbf{F}}^{2} \max\{\|W\|_{\mathbf{F}}^{2}, R^{2}/4^{n}\} \sup_{\epsilon > 0} M_{l \to L}^{2}(W,\epsilon)}{n}} \log \frac{1}{\alpha} + (1-\alpha) \sqrt{\sum_{l=1}^{L} \log(d_{l-1}n)}, \end{split}$$

where $C_3 > 0$ is an absolute constant.

 Step 2: Bounding the Rest Integral. We then prove $\int_0^{\alpha} \sqrt{d_R(W, \epsilon)} d\epsilon$. Again, by the basic inequality $\log x \le \log(1+x) \le x$ for x > 0, we have

$$\begin{split} &\sum_{k=1}^{r_{\text{eff}}[W,l]} \log \left(\frac{8C_{2}^{2}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2},R^{2}/4^{n}\}LM_{l \to L}^{2}(W,\epsilon)}{n\epsilon^{2}} \right) \\ &\leq \sum_{k=1}^{d_{l-1}} \log \left(\frac{8C_{2}^{2}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2},R^{2}/4^{n}\}LM_{l \to L}^{2}(W,\epsilon)}{n\epsilon^{2}} \right) \\ &= \sum_{k=1}^{d_{l-1}} \log \left(\frac{8C_{2}^{2}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2},R^{2}/4^{n}\}LM_{l \to L}^{2}(W,\epsilon)}{n\alpha^{2}} \right) + d_{l-1} \log \frac{\alpha^{2}}{\epsilon^{2}} \\ &\leq \frac{8C_{2}^{2}\sum_{k=1}^{d_{l-1}}\lambda_{k}(F_{l-1}F_{l-1}^{\top}) \cdot \max\{||W||_{F}^{2},R^{2}/4^{n}\}LM_{l \to L}^{2}(W,\epsilon)}{n\alpha^{2}} + d_{l-1} \log \frac{\alpha^{2}}{\epsilon^{2}} \\ &= \frac{8C_{2}^{2}||F_{l-1}(W,X)||_{\mathbf{F}}^{2} \cdot \max\{||W||_{F}^{2},R^{2}/4^{n}\}LM_{l \to L}^{2}(W,\epsilon)}{n\alpha^{2}} + d_{l-1} \log \frac{\alpha^{2}}{\epsilon^{2}}. \end{split} \tag{84}$$

Taking (84) to the integral $\int_0^\alpha \sqrt{d_R(W,\epsilon)} d\epsilon,$ we have

$$\begin{split} & \int_{0}^{\alpha} \sqrt{d_{R}(W,\epsilon)} d\epsilon \\ \leq & 2\sqrt{2}C_{2} \int_{0}^{\alpha} \sqrt{\sum_{l=1}^{L} (d_{l} + d_{l-1}) \frac{\|F_{l-1}\|_{\mathbf{F}}^{2} \|W\|_{\mathbf{F}}^{2} L M_{l \to L}^{2}(W,\epsilon)}{n\alpha^{2}}} d\epsilon + \int_{0}^{\alpha} \sqrt{\sum_{l=1}^{L} (d_{l} + d_{l-1}) d_{l-1} \log \frac{\alpha^{2}}{\epsilon^{2}}} d\epsilon \\ \leq & C_{4} \left(\sqrt{\frac{\sum_{l=1}^{L} (d_{l} + d_{l-1}) \|F_{l-1}\|_{\mathbf{F}}^{2} \max\{\|W\|_{\mathbf{F}}^{2}, R^{2} / 4^{n}\} L \sup_{\epsilon > 0} M_{l \to L}^{2}(W,\epsilon)}{n}} + \alpha \sqrt{\sum_{l=1}^{L} (d_{l} + d_{l-1}) d_{l-1}} \right), \end{split}$$

where the second inequality holds by calculating the integral $\int_0^\alpha \sqrt{\log\left(\frac{\alpha^2}{\epsilon^2}\right)}d\epsilon = \alpha\sqrt{\frac{\pi}{2}}$, and $C_4>0$ is an absolute constant. Taking $\alpha=\frac{1}{n^5}$, the high-order term $\alpha\sqrt{\sum_{l=1}^L(d_l+d_{l-1})d_{l-1}}$ will be $\frac{\sqrt{\sum_{l=1}^L(d_l+d_{l-1})d_{l-1}}}{n^5}$ and is ignorable.

Step 3: Combining the Two Integrals. Combining Step 1 and Step 2, we get the full Riemannian Dimension integral upper bound

$$\frac{1}{\sqrt{n}} \int_0^\infty \sqrt{d_R(W, \epsilon)} d\epsilon \le O\left(\frac{\sqrt{\sum_{l=1}^L (d_l + d_{l-1}) L \|F_{l-1}\|_F^2 \|W\|_F^2 \sup_{\epsilon > 0} M_{l \to L}^2(W, \epsilon)}}{n} + \sqrt{\frac{\sum_{l=1}^L \log(d_{l-1}n)}{n}}\right),$$

where O hides multiplicative absolute constants and two ignorable high-order terms: $\frac{\beta\sqrt{\sum_{l=1}^L(d_l+d_{l-1})d_{l-1}}}{n^{5.5}} \text{ and } \frac{\beta\sqrt{\sum_{l=1}^L(d_l+d_{l-1})L\|F_{l-1}\|_{\mathbf{F}}^2R^2\sup_{\epsilon>0}M_{l\to L^2(W,\epsilon)}}}{n2^n}.$

Put this bound into Theorem 4 (or (73) in its proof), we have with probability at least $1-\delta$, uniformly over all $w \in \mathcal{W}$,

$$(\mathbb{P} - \mathbb{P}_{n})\ell(f(W, x), y) \\
\leq O\left(\frac{\beta\sqrt{\sum_{l=1}^{L}(d_{l} + d_{l-1})L\|F_{l-1}\|_{\mathbf{F}}^{2}\|W\|_{\mathbf{F}}^{2}\sup_{\epsilon > 0}M_{l \to L}^{2}(W, \epsilon)}}{n} + \sqrt{\frac{\beta^{2}\sum_{l=1}^{L}\log(d_{l-1}n) + \log\frac{\log(2n)}{\delta}}{n}}\right), \tag{85}$$

where O hides multiplicative absolute constants and two ignorable high-order terms: $\frac{\beta\sqrt{\sum_{l=1}^L(d_l+d_{l-1})d_{l-1}}}{n^{5.5}} \quad \text{and} \quad \frac{\beta\sqrt{\sum_{l=1}^L(d_l+d_{l-1})L\|F_{l-1}\|_{\mathbf{F}}^2R^2\sup_{\epsilon>0}M_{l\to L^2(W,\epsilon)}}}{n2^n} \quad \text{Note that here} \quad F_{l-1}(W;X) \in R^{d_{l-1}\times n} \quad \text{contains } n \text{ features vectors in dimension } d_{l-1} \text{ so its Frobenius norm}$

 $||F_{l-1}||_{\mathsf{F}}$ scales with \sqrt{n} with respect to sample size; and $\sup_{\epsilon>0} M_{l\to L}(W,\epsilon)$ is the "one-point" Lipchitz constant at W in the sense that

$$||F_L(F_l(W',X),\{W_i'\}_{i=l+1}^L) - F_L(F_l(W,X),\{W_i'\}_{i=l+1}^L)||_{\mathbf{F}}$$

$$\leq \left(\sup_{\epsilon} M_{l\to L}(W,\epsilon)\right) ||F_l(W',X) - F_l(W,X)||_{\mathbf{F}}, \quad \forall W' \in \mathcal{W}.$$

This concludes the first generalization bound in Corollary 1.

Step 4: Prove the Second Generalization Bound. Now we continue to show that the bound in Corollary 1 is strictly better than the spectrally normalized bound. To see this, as we presented under Corollary 1, we have

$$||F_{l-1}(W,X)||_{\mathbf{F}}$$

$$= ||\sigma_{l-1}(W_{l-1} \cdots W_2 \sigma_1(W_1 X))||_{\mathbf{F}}$$

$$\leq \Pi_{i < l} ||W_i||_{op} \cdot ||X||_{\mathbf{F}}, \tag{86}$$

by the property of spectral norm ($\|AB\|_{\mathbf{F}} \leq \|A\|_{\mathrm{op}} \|B\|_{\mathbf{F}}$)), and the fact that all activation functions are 1-Lipchitz in column.

In the meanwhile, we know that

$$\left(\sup_{\epsilon} M_{l\to L}(W,\epsilon)\right) \leq \sup_{\epsilon} \Pi_{i>l} \|W_i'\|_{\text{op}},$$

again by the property of spectral norm ($\|AB\|_{\mathbf{F}} \leq \|A\|_{\mathrm{op}} \|B\|_{\mathbf{F}}$)) and the fact that all activation functions are 1-Lipchitz in column. This results in

$$\sup_{\epsilon} \Pi_{i>l} \|W_i'\|_{\text{op}} \le \sup_{w \in \mathcal{W}} \Pi_{i>l} \|W_i\|_{\text{op}}. \tag{87}$$

Combining (86) and (87) together with (85), we have that for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, uniformly over every $w \in \mathcal{W}$, we have

$$(\mathbb{P} - \mathbb{P}_{n})\ell(f(W, x), y) \\
\leq O\left(\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2} \cdot \sum_{l=1}^{L} (d_{l} + d_{l-1})\Pi_{i < l}\|W_{i}\|_{op}^{2} \sup_{w \in \mathcal{W}} \Pi_{i > l}\|W_{i}\|_{op}^{2}}{n} + \sqrt{\frac{\beta^{2} \sum_{l=1}^{L} \log(d_{l-1}n) + \log \frac{\log(2n)}{\delta}}{n}}\right), \tag{88}$$

where O hides multiplicative absolute constants and two ignorable high-order terms: $\frac{\beta\sqrt{\sum_{l=1}^{L}(d_l+d_{l-1})d_{l-1}}}{n^{5.5}} \text{ and } \frac{\beta\sqrt{\sum_{l=1}^{L}(d_l+d_{l-1})L\|F_{l-1}\|_{\mathbf{F}}^2R^2\sup_{\epsilon>0}M_{l\to L^2(W,\epsilon)}}}{n2^n}.$

The next step is to use Lemma 8 (the "uniform pointwise convergence" principle Xu & Zeevi (2025)) to give a conversion from the uniform convergence to the pointwise convergence. Denote the functional $T_l: \mathcal{W} \to (0, R_l]$ is defined by

$$T_l(W) = \prod_{i \neq l} ||W_i||_{\text{op}}^2.$$

Since $\sum_{i\neq l} \|W_i\|_{\mathsf{F}}^2 \leq \|W\|_{\mathsf{F}}^2 \leq R^2$, we have $T_l(W) = \prod_{i\neq l} \|W_i\| \leq (R/\sqrt{L-1})^{L-1}$ according to the AM-GM inequality. The bound in (88) implies that for any $l=1,\cdots,L$, $\forall t_l \in (0,(R/\sqrt{L-1})^{L-1})]$, with probability at least $1-\delta$,

$$\sup_{W:T_{l}(W)\leq t_{l},\forall l\in[L]} (\mathbb{P}-\mathbb{P}_{n})\ell(W;z)$$

$$\leq O\left(\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2}\cdot\sum_{l=1}^{L}(d_{l}+d_{l-1})t_{l}^{2}}}{n}+\sqrt{\frac{\beta^{2}\sum_{l=1}^{L}\log(d_{l-1}n)+\log\frac{\log(2n)}{\delta}}{n}}\right). \quad (89)$$

The inequality (89) precisely match the condition in Lemma 8 for each functional $d_l(w)$. Thus applying Lemma 8 by L times, with the smallest radius r_0 in the statement of Lemma 8 chosen to be $r_0 = (R/\sqrt{L-1})^{L-1})/\max\{R,2\}^n$, we have that for any $\delta \in (0,1)$, with probability at least $1-\delta$, uniformly over every $w \in \mathcal{W}$,

$$\begin{aligned}
&(\mathbb{P} - \mathbb{P}_{n})\ell(f(W, x), y) \\
&\leq O\left(\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2}\sum_{l=1}^{L}(d_{l} + d_{l-1})\max\{4T_{l}^{2}(W), \frac{(R/\sqrt{L-1})^{2L-2}}{\max\{R, 2\}^{2n}}\}}{n}\right) \\
&+ \sqrt{\frac{\beta^{2}\sum_{l=1}^{L}\log(d_{l-1}n) + L\log\frac{n\log\max\{R, 2\}}{\delta}}{n}}\right) \\
&= O\left(\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2} \cdot \sum_{l=1}^{L}(d_{l} + d_{l-1})\prod_{i \neq l}\|W_{i}\|_{\mathsf{op}}^{2}}}{n} + \sqrt{\frac{\beta^{2}\sum_{l=1}^{L}\log(d_{l-1}n) + L\log\frac{n\log\max\{R, 2\}}{\delta}}{n}}\right)} \\
&= O\left(\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2} \cdot \sum_{l=1}^{L}(d_{l} + d_{l-1})\prod_{i \neq l}\|W_{i}\|_{\mathsf{op}}^{2}}}{n}\right) \\
&= O\left(\frac{\beta\sqrt{L\|W\|_{\mathsf{F}}^{2}\|X\|_{\mathsf{F}}^{2} \cdot \sum_{l=1}^{L}(d_{l} + d_{l-1})\prod_{i \neq l}\|W_{i}\|_{\mathsf{op}}^{2}}}{n}\right)
\end{aligned}$$
(90)

where
$$O$$
 hides multiplicative absolute constants and three ignorable high-order terms:
$$\frac{\beta\sqrt{\sum_{l=1}^{L}(d_{l}+d_{l-1})d_{l-1}}}{n^{5.5}}, \frac{\beta\sqrt{\sum_{l=1}^{L}(d_{l}+d_{l-1})L\|F_{l-1}\|_{\mathbf{F}}^{2}R^{2}\sup_{\epsilon>0}M_{l\to L^{2}(W,\epsilon)}}}{n^{2^{n}}} \quad \text{and} \quad \frac{\beta\sqrt{L\|W\|_{\mathbf{F}}^{2}\|X\|_{\mathbf{F}}^{2}\sum_{l=1}^{L}(d_{l}+d_{l-1})(R/\sqrt{L-1})^{L-1}}}{n\max\{R,2\}^{n}}.$$

Now we see from (86) and (87) that the derived norm-constraint bound (85) implies the spectrally normalized bound (90). This completes the proof.

Discussion of Corollary 1. We proceed in three paragraphs of discussion. First, we show that the Riemannian–Dimension bound in Theorem 4 is *exponentially* tighter than the spectrally normalized bound in (80). Second, we offer a metric–tensor interpretation that clarifies the source of this improvement. Finally, we position (80) relative to the most representative SNB bounds in the existing literature.

I: Why the improvement is exponential. Empirically one observes

$$||F_{l-1}||_{\mathbf{F}} \ll \Pi_{i < l} ||W_i||_{\text{op}} ||X||_{\mathbf{F}}, \qquad M_{l \to L}(W, \varepsilon) \leq \sup_{W' \in B_{\varrho_n}(W, \varepsilon)} \Pi_{i > l} ||W_i'||_{\text{op}},$$

so Theorem 4 is *exponentially* tighter than (80). Therefore, Theorem 4 improves on Corollary 1 by an exponential factor.

II: Metric tensor interpretation. For understand the improvement deeper, we highlight that the spectral norm bound (80) can be equivalently viewed as replacing the metric tensor $G_{\rm NP}$ (6) used in Theorem 4 by the diagonal metric tensor

$$G_{\mathrm{SNB}}(W) = \mathrm{blockdiag}\Big(\dots, L\sup_{W' \in \mathcal{W}} \Pi_{k \neq l} \|W_k'\|_{\mathrm{op}} \|X\|_{\mathsf{F}}^2 \otimes I_{d_l \times d_{l-1}}, \dots\Big),$$

which is a far coarser relaxation that completely discards the learned feature $F_l(W, X)$.

III: Relation to existing spectrally normalised bounds. The bound in (80) is structurally close to the classical SNB results of Bartlett et al. (2017) and Neyshabur et al. (2018); the three bounds differ only in the *global ball* used to constraint the hypothesis class.

- (a) Our bound (80) controls *all* layers simultaneously via the global Frobenius norm $||W||_{\mathbf{F}}$, hence the factor $||W||_{\mathbf{F}}$ in the numerator.
- (b) Neyshabur et al. (2018) bounds each layer l separately by its Frobenius norm $||W_l||$. Strengthening their argument with Dudley's entropy integral (one-shot optimization in the

original paper) gives

$$(\mathbb{P} - \mathbb{P}_n) \, \ell \big(f(W, x), y \big) \, \leq \, \tilde{O} \Big(\frac{\beta \, \|X\|_{\mathsf{F}} \, \sqrt{\sum_{l=1}^L L^2(d_l + d_{l-1}) \, \|W_l\|_{\mathsf{F}}^2 \, \Pi_{i \neq l} \|W_i\|_{\mathsf{op}}^2}}{n} + \sqrt{\frac{\log \frac{1}{\delta}}{n}} \Big).$$

Neither (80) nor (91) strictly dominates the other, since factors of the form $(\sum_l a_l)(\sum_l b_l)$ in (80) vs. factors of the form $L\sum_l a_l b_l$ in (91) can swap their relative order.

(c) Bartlett et al. (2017) replaces each Frobenius norm by the $\|\cdot\|_{2,1}$ norm, obtaining the tighter

$$(\mathbb{P} - \mathbb{P}_n) \, \ell \big(f(W, x), y \big) \, \leq \, \tilde{O} \Big(\frac{\beta \, \|X\|_{\mathsf{F}} \, \big[\sum_l \|W_l\|_{2, 1}^{2/3} \, \sum_l \big(\Pi_{i \neq l} \|W_i\|_{\mathsf{op}}^2 \big)^{2/3} \big]^{3/2}}{n} + \sqrt{\frac{\log \frac{1}{\delta}}{n}} \Big), \tag{92}$$

which improves on (80) and (91) thanks to the sharper 2,1 norm. Extending our Riemannian-dimension analysis to the 2,1 norm setting is an interesting direction for future work.

(d) Size-independent SNB bounds (pioneered by Golowich et al. (2020)) remove all depth/width dependence at the price of a worse scaling in n; incorporating their technique is left for future research.

In any case, (80) is a representative SNB bound, and the key message in this subsection is that our Riemannian-Dimension result in Theorem 4 is *exponentially* sharper than (80).