
Discrepancies are Virtue: Weak-to-Strong Generalization through Lens of Intrinsic Dimension

Yijun Dong¹ Yicheng Li¹ Yunai Li² Jason D. Lee³ Qi Lei¹

Abstract

Weak-to-strong (W2S) generalization is a type of finetuning (FT) where a strong (large) student model is trained on pseudo-labels generated by a weak teacher. Surprisingly, W2S FT often outperforms the weak teacher. We seek to understand this phenomenon through the observation that FT often occurs in intrinsically low-dimensional spaces. Leveraging the low intrinsic dimensionality of FT, we analyze W2S in the ridgeless regression setting from a variance reduction perspective. For a strong student-weak teacher pair with sufficiently expressive low-dimensional feature subspaces $\mathcal{V}_s, \mathcal{V}_w$, we provide an exact characterization of the variance that dominates the generalization error of W2S. This unveils a virtue of discrepancy between the strong and weak models in W2S: the variance of the weak teacher is inherited by the strong student in $\mathcal{V}_s \cap \mathcal{V}_w$, while reduced by a factor of $\dim(\mathcal{V}_s)/N$ in the subspace of discrepancy $\mathcal{V}_w \setminus \mathcal{V}_s$ with N pseudo-labels for W2S. Our analysis further casts light on the sample complexities and the scaling of performance gap recovery in W2S. The analysis is supported by experiments on synthetic regression problems, as well as real vision and NLP tasks.

1. Introduction

As the capabilities of modern machine learning models grow and exceed human performance in many domains, an emerging problem is whether it would be possible to align the strong superhuman models with weaker supervisors such as human feedback. The weak-to-strong (W2S) framework introduced in (Burns et al., 2024) is a feasible analogy for this problem, inquiring how much capacity of a strong student model can be evoked under the supervision

of a weak teacher model. W2S is related to various learning paradigms like co-training (Blum & Mitchell, 1998), self-training (Scudder, 1965), knowledge distillation (Hinton, 2015), and self-distillation (Zhang et al., 2019; 2021), yet being critically dissimilar.

Formalizing the *discrepancy* between the student and the teacher in their *model capacities* is essential for understanding W2S. Most existing theories for W2S treat model capacity as an absolute notion with respect to the downstream task, e.g. the weak teacher lacks the robustness to perturbation (Lang et al., 2024; Shin et al., 2024) or the ability to fit the target function (Ildiz et al., 2024; Wu & Sahai, 2024). Nevertheless, empirical observations suggest W2S models also surpass weak models’ performance on less challenging tasks (Burns et al., 2024), where the weak teacher has sufficient capacity to achieve good performance. This gap of understanding motivates some natural questions:

Why W2S happens when both the teacher and student have sufficient capacities for the downstream task?

What affects W2S generalization beyond the absolute notion of model capacity?

To answer the above questions, we analyze W2S generalization through the lens of intrinsic dimension beyond the absolute notion of model capacity. We develop a theoretical framework that incorporates student-teacher correlation, providing a more nuanced explanation of when and why W2S model surpasses the weak teacher’s performance.

Our framework is built on two inspiring observations on finetuning (FT): (i) FT tends to fall in the kernel regime (Jacot et al., 2018; Wei et al., 2022; Malladi et al., 2023); and (ii) for a downstream task, relevant features in a stronger pretrained model tend to concentrate in a subspace of lower dimension, known as the *intrinsic dimension*, even when FT is highly overparametrized (Aghajanyan et al., 2021). Leveraging these properties, we cast FT as a ridgeless regression problem over subgaussian features. In particular, we consider two subspaces $\mathcal{V}_s, \mathcal{V}_w \subset \mathbb{R}^d$ of low dimensions $d_s, d_w \ll d$ that encapsulate relevant features in the strong student and weak teacher, respectively. The “absolute” model capacities are measured from two aspects: (i) the intrinsic dimensions d_s, d_w that quantify the representation

¹New York University ²Shanghai Jiaotong University
³Princeton University. Correspondence to: Yijun Dong <yid1319@nyu.edu>, Qi Lei <q1518@nyu.edu>.

Proceedings of the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

“complexity” and (ii) the approximation errors $\rho_s < \rho_w$ that quantify the representation “accuracy” of the strong and weak models, respectively. In addition, the student-teacher correlation is measured by alignment between the strong and weak feature subspaces through their canonical angles (see Appendix D), $d_{s \wedge w} = \sum \cos(\angle(\mathcal{V}_s, \mathcal{V}_w))$ such that $d_{s \wedge w} \in [0, \min\{d_s, d_w\}]$.

This framework reveals the roles of low intrinsic dimensions and student-teacher correlation in W2S. Decomposing the W2S generalization error into variance and bias, the bias is due to the approximation errors, ρ_s, ρ_w , which are low when both student and teacher have sufficient capabilities; whereas the variance comes from noise in the labeled samples for finetuning the weak teacher. When finetuning the strong student with $N \gtrsim d_s$ pseudo-labels generated by a weak teacher supervisedly finetuned with $n \gtrsim d_w$ noisy labels, the variance of W2S is proportional to:

$$\underbrace{\frac{d_{s \wedge w}}{n}}_{\text{Var. in } \mathcal{V}_s \cap \mathcal{V}_w} + \underbrace{\frac{d_s}{N}}_{\text{W2S}} \underbrace{\frac{d_w - d_{s \wedge w}}{n}}_{\text{Var. in } \mathcal{V}_w \setminus \mathcal{V}_s}.$$

Specifically, the student mimics variance of the weak teacher in the overlapped feature subspace $\mathcal{V}_s \cap \mathcal{V}_w$ but reduces the variance by a factor of d_s/N in the discrepancy between \mathcal{V}_w and \mathcal{V}_s . Compared to the weak teacher variance that scales as d_w/n , W2S happens (*i.e.* the student outperforms its weak teacher) with sufficient sample sizes n, N when: (i) the strong student has a lower intrinsic dimension, $d_s < d_w$ (as empirically observed in (Aghajanyan et al., 2021) on NLP tasks), or (ii) the student-teacher correlation is low, $d_{s \wedge w} < d_w$. This unveils the benefit of discrepancy between the teacher and student features for W2S:

In the variance-dominated regime, W2S comes from variance reduction in the discrepancy of weak teacher features from strong student features.

To provide intuitions for such variance reduction, let’s consider \mathcal{V}_s and \mathcal{V}_w with large discrepancy as two distinct aspects of a downstream task that both provide sufficient information. For example, to classify the brand of a car in an image, one can use either the simple information in the logo (strong features \mathcal{V}_s with a lower intrinsic dimension d_s) or the complex information in the design (weak features \mathcal{V}_w with a higher intrinsic dimension d_w). In a high-dimensional feature space, \mathcal{V}_s and \mathcal{V}_w that encode irrelevant information are likely almost orthogonal, leading to a small $d_{s \wedge w}$. Then, the error of weak teacher induced by noise in the n labeled samples is only correlated to the design features in \mathcal{V}_w but almost independent of the logo features in \mathcal{V}_s . Therefore, the error in weak supervision can be viewed as independent label noise for the student. With an intrinsic dimension of d_s , the generalization error of student induced by such independent noise vanishes at a rate of $O(d_s/N)$.

Our main contributions are summarized as follows:

- We introduce a theoretical framework for W2S based on the low intrinsic dimensions of FT, where we characterize model capacities from three aspects: approximation errors for “accuracy”, intrinsic dimensions for “complexity”, and student-teacher correlation for “alignment” (Section 2).
- We provide a generalization analysis for W2S with an exact characterization of the variance under a Gaussian feature assumption, unveiling the virtue of discrepancy between the student and teacher in W2S (Section 3.1).
- We investigate the relative W2S performance in terms of performance gap recovery (PGR) (Burns et al., 2024) and outperforming ratio (OPR) compared to the strong baseline model supervisedly finetuned with n labels. A case study provides insights into the scaling of PGR and OPR with respect to the sample sizes n, N and sample complexities in W2S (Section 3.2).

1.1. Related works

In this section, we review literature directly related to W2S and intrinsic dimension, while deferring detailed discussions on other related topics to Appendix A.

W2S alignment: emergence & growing influence. W2S generalization was first introduced by (Burns et al., 2024), offering a promising avenue for aligning superhuman models. A rapidly expanding body of work has empirically validated this phenomenon across diverse tasks in vision and language models since then. Guo et al. (2024); Liu & Alahi (2024) propose loss functions and multi-teacher algorithms. Guo & Yang (2024); Yang et al. (2024b) refine training data to improve W2S alignment, while Li et al. (2024); Sun et al. (2024) use weak models for data filtering and reranking. In contrast, Yang et al. (2024a) highlight the issue of W2S deception, where strong models superficially align with weak teachers but fail in new or conflicting cases. This calls for theoretical understanding of the mechanism behind W2S generalization and better strategies to mitigate misalignment. Coinciding with our theoretical findings, the strong negative correlation between model similarity and W2S performance is empirically observed in the concurrent work (Goel et al., 2025) in extensive experiments.

Theoretical perspectives on W2S generalization. Existing theories on W2S interpret the difference between strong and weak models in terms of the quality of their representations (from the bias perspective in our context). Lang et al. (2024) study W2S in classification through the lens of neighborhood expansion (Wei et al., 2020; Cai et al., 2021) where model capacity is interpreted as the robustness to perturbation. Within this framework, Shin et al. (2024) highlights the importance of data selection in W2S while

proposing metrics and algorithms for data selection in W2S. In the same classification setting, Somerstep et al. (2024) takes a transfer learning perspective and highlights the limitation of naive FT in W2S. Wu & Sahai (2024) take a benign overfitting (Bartlett et al., 2020; Muthukumar et al., 2021) perspective and show the asymptotic transition between W2S generalization and random guessing. For regression tasks, Charikar et al. (2024) reveals the connection between W2S gain and misfit error of the strong student on weak pseudo-labels. Ildiz et al. (2024) treats W2S as a special case of knowledge distillation, showing its limitation in terms of improving the data scaling law (Spigler et al., 2020; Bahri et al., 2024). We consider a similar ridgeless regression setting as Ildiz et al. (2024) but from a fundamentally different aspect – variance reduction. This offers a fresh take on the roles of intrinsic dimension and student-teacher correlation in W2S. Parallel to this work, Medvedev et al. (2025); Yao et al. (2025) analyze W2S in the context of early stopping and loss functions, respectively.

Intrinsic dimension. There has been prevailing empirical and theoretical evidence that natural high-dimensional systems often exhibit low-dimensional structures (Udell & Townsend, 2019). The concept of intrinsic dimension has been widely studied in manifold learning (Tenenbaum et al., 2000), dimensionality reduction (Van der Maaten & Hinton, 2008), and representation learning (Bengio et al., 2013). In the context of neural network training, Li et al. (2018) propose a method to measure the intrinsic dimension of the objective landscape based on the Johnson-Lindenstrauss-type transforms (Johnson, 1984). This offers a structural perspective on task complexity, which is largely absent from prior W2S studies. Aghajanyan et al. (2021) investigate the intrinsic dimensions of FT, showing that FT over large models usually has surprisingly low intrinsic dimensions, while good pretraining tends to reduce the intrinsic dimension. Our work extends these insights by linking the intrinsic dimension to W2S, decomposing generalization error into bias and variance, and building upon findings from Yang et al. (2020); Amari et al. (2020) on variance-dominated risks in learning from noisy labels.

1.2. Notations

Given any $n \in \mathbb{Z}_+$, we denote $[n] = \{1, \dots, n\}$. Let \mathbf{e}_n be the n -th canonical basis of conformable dimension; \mathbf{I}_n is the $n \times n$ identity matrix; and $\mathbf{0}_n, \mathbf{1}_n \in \mathbb{R}^n$ are vectors with all zeroes and ones. For any distribution p and $n \in \mathbb{Z}_+$, let $p^n \triangleq \bigotimes_{i=1}^n p$ as the n -fold product distribution of p , sampling which yields n i.i.d. samples from p . For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, let \mathbf{A}^\dagger be the Moore-Penrose pseudoinverse. We adapt the standard asymptotic notations: for any functions $f, g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, we write $f = O(g)$ or $f \lesssim g$ if there exists some constant $C > 0$ such that $f(x) \leq Cg(x)$ for all $x \in \mathbb{R}_+$; $f = \Omega(g)$ or $f \gtrsim g$ if $g = O(f)$; $f \asymp g$

if $f = O(g)$ and $f = \Omega(g)$. Also, we denote $f = o(g)$ or $f/g = o_x(1)$ if $\lim_{x \rightarrow \infty} f(x)/g(x) = 0$.

2. Problem setup

In this section, we cast FT as a ridgeless regression problem. The setup is introduced in three parts: model capacity, FT algorithms, and metrics for W2S performance.

Consider the problem of learning an unknown data distribution $\mathcal{D}(f_*) : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ (where \mathcal{X} is a set and $\mathcal{Y} \subseteq \mathbb{R}$) associated with a downstream task characterized by an unknown ground truth function $f_* : \mathcal{X} \rightarrow \mathbb{R}$. Every sample $(\mathbf{x}, y) \sim \mathcal{D}(f_*)$ satisfies $y = f_*(\mathbf{x}) + z$ where $z \sim \mathcal{N}(0, \sigma^2)$ is an independent Gaussian label noise. Let $\mathcal{D} : \mathcal{X} \rightarrow [0, 1]$ be the marginal distribution over \mathcal{X} . We assume that f_* is bounded: $|f_*(\mathbf{x})| \leq 1$ for $\mathbf{x} \sim \mathcal{D}$ almost surely (under normalization without loss of generality).

For any $f : \mathcal{X} \rightarrow \mathbb{R}$ learned with samples from $\mathcal{D}(f_*)$, we measure the generalization error via excess risk:

$$\mathbf{ER}(f) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\mathbb{E}_f [(f(\mathbf{x}) - f_*(\mathbf{x}))^2]], \quad (1)$$

where randomness in f comes from its training samples. Notice that $\mathbf{ER}(f)$ can be decomposed into variance and bias, $\mathbf{ER}(f) = \mathbf{Var}(f) + \mathbf{Bias}(f)$, where

$$\mathbf{Var}(f) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\mathbb{E}_f [(f(\mathbf{x}) - \mathbb{E}_f[f(\mathbf{x})])^2]],$$

$$\mathbf{Bias}(f) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\mathbb{E}_f[f(\mathbf{x})] - f_*(\mathbf{x}))^2].$$

2.1. Measures for model capacity

Model capacity is a key notion in W2S that distinguishes the weak and strong models. Intuitively, a stronger model is capable of representing a downstream task $\mathcal{D}(f_*)$ more accurately and efficiently. We formalize such ‘‘accuracy’’ and ‘‘complexity’’ through the notions of intrinsic dimensions and FT approximation errors, as introduced below.

Consider two pretrained models, a weak model ϕ_w and a strong model ϕ_s , that output features $\mathcal{X} \rightarrow \mathbb{R}^d$:

Assumption 2.1 (Sub-gaussian features). For $\mathbf{x} \sim \mathcal{D}$, assume both $\phi_w(\mathbf{x})$ and $\phi_s(\mathbf{x})$ are zero-mean sub-gaussian random vectors with $\mathbb{E}[\phi_w(\mathbf{x})] = \mathbb{E}[\phi_s(\mathbf{x})] = \mathbf{0}_d$, and $\mathbb{E}[\phi_w(\mathbf{x})\phi_w(\mathbf{x})^\top] = \Sigma_w$, $\mathbb{E}[\phi_s(\mathbf{x})\phi_s(\mathbf{x})^\top] = \Sigma_s$.

Approximation errors measure the model capacity from the ‘‘accuracy’’ perspective: how accurately can the downstream task $\mathcal{D}(f_*)$ be represented by the pretrained features of ϕ_s and ϕ_w over the population.

Definition 2.2 (FT approximation error). Given $\mathcal{D}(f_*)$, let the FT approximation errors of ϕ_s and ϕ_w be

$$\rho_s = \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_s(\mathbf{x})^\top \boldsymbol{\theta} - f_*(\mathbf{x}))^2],$$

$$\rho_w = \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_w(\mathbf{x})^\top \boldsymbol{\theta} - f_*(\mathbf{x}))^2],$$

such that $\rho_s, \rho_w \in [0, 1]$ (given $\Pr_{\mathbf{x} \sim \mathcal{D}}[|f_*(\mathbf{x})| \leq 1] = 1$ by assumption). We assume both ρ_s and ρ_w are small compared to label noise: $\rho_s + \rho_w \ll \sigma^2$; while the stronger model ϕ_s has a lower FT approximation error: $\rho_s < \rho_w$. Further, with respect to a sample size n , let $\rho_s(n) = \mathbb{E}_{\mathbf{X} \sim \mathcal{D}^n}[\|\phi_s(\mathbf{X})\phi_s(\mathbf{X})^\top f_*(\mathbf{X}) - f_*(\mathbf{X})\|_2^2]$ and $\rho_w(n) = \mathbb{E}_{\mathbf{X} \sim \mathcal{D}^n}[\|\phi_w(\mathbf{X})\phi_w(\mathbf{X})^\top f_*(\mathbf{X}) - f_*(\mathbf{X})\|_2^2]$, where $\phi_s(\mathbf{X})$ and $\phi_w(\mathbf{X})$ are $n \times d$ feature matrices; and $f_*(\mathbf{X}) \in \mathbb{R}^n$ is a vector of the noiseless ground truth labels.

Notice that FT approximation error is different from approximation error of the full model. Precisely, Definition 2.2 quantifies the approximation error of *finetuning the pretrained model*, whose dynamics (Wei et al., 2022; Malladi et al., 2023) fall in the kernel regime (Jacot et al., 2018). Since feature learning is limited in the kernel regime (Woodworth et al., 2020), a low FT approximation error requires the pretrained features ϕ_s and ϕ_w to provide an expressive set of features for the downstream task $\mathcal{D}(f_*)$.

In addition to “accuracy”, a strong model ought to be able to represent a downstream task concisely. We quantify such “complexity” through intrinsic dimension – the minimum dimension of a feature subspace that can represent the downstream task $\mathcal{D}(f_*)$ accurately. In light of the ubiquitous observations on low intrinsic dimensions of FT (Aghajanyan et al., 2021), we introduce a common assumption for FT (Xia et al., 2024; Dong et al., 2024c) that the pretrained features of ϕ_s, ϕ_w are concentrated in low-dimensional subspaces, as formalized below.

Definition 2.3 (Intrinsic dimensions). Let $d_s = \text{rank}(\Sigma_s)$ and $d_w = \text{rank}(\Sigma_w)$ be the **intrinsic dimensions** of ϕ_s and ϕ_w . Assume low intrinsic dimensions¹: $d_s, d_w \ll d$.

Moreover, Aghajanyan et al. (2021) observed that the stronger pretrained models tend to have lower intrinsic dimensions, *i.e.* we often have $d_s < d_w$ in practice.

Beyond the absolute notion of model capacity in terms of intrinsic dimensions and FT approximation errors, we introduce a relative measure for the similarity between weak and strong models – the correlation dimension characterizing the overlap between feature subspaces of ϕ_s and ϕ_w .

Definition 2.4 (Correlation dimension). Consider spectral decompositions $\Sigma_s = \mathbf{V}_s \Lambda_s \mathbf{V}_s^\top$ and $\Sigma_w = \mathbf{V}_w \Lambda_w \mathbf{V}_w^\top$, where $\Lambda_s \in \mathbb{R}^{d_s \times d_s}$ and $\Lambda_w \in \mathbb{R}^{d_w \times d_w}$ are diagonal matrices with positive eigenvalues in decreasing order; while

¹ In practice, Σ_s and Σ_w usually admit fast-decaying eigenvalues, but not exactly low-rank. In this more realistic case, ridge regression with suitable choices of regularization hyperparameters intuitively performs “soft” truncation of the small singular values, effectively leading to low intrinsic dimensions $d_s, d_w \ll d$. For conciseness of the main message, we focus on the ideal case of exactly low-rank Σ_s and Σ_w in the main text, while deferring the ridge regression analysis for general Σ_s and Σ_w to Appendix C.

$\mathbf{V}_s \in \mathbb{R}^{d \times d_s}$ and $\mathbf{V}_w \in \mathbb{R}^{d \times d_w}$ consist of the corresponding orthonormal eigenvectors. Let $d_{s \wedge w} = \|\mathbf{V}_s^\top \mathbf{V}_w\|_F^2$ be the **correlation dimension** between ϕ_s and ϕ_w such that $0 \leq d_{s \wedge w} \leq \min\{d_s, d_w\}$.

Remark 2.5 (Extension to general FT). While we focus on learning $\mathcal{D}(f_*)$ via linear probing over ϕ_w and ϕ_s , since the finetuning dynamics fall approximately in the kernel regime (Wei et al., 2022; Malladi et al., 2023), the linear probing analysis naturally extends to general FT. Precisely, let $f_w(\cdot|\mathbf{0}_d) : \mathcal{X} \rightarrow \mathbb{R}$ and $f_s(\cdot|\mathbf{0}_d) : \mathcal{X} \rightarrow \mathbb{R}$ be the pretrained weak and strong models, where d is the number of finetunable parameters. By denoting $\phi_w(\mathbf{x}) = \nabla_{\theta} f_w(\mathbf{x}|\mathbf{0}_d)$ and $\phi_s(\mathbf{x}) = \nabla_{\theta} f_s(\mathbf{x}|\mathbf{0}_d)$, the general FT process effectively reduces to linear probing over ϕ_w and ϕ_s .

2.2. W2S and supervised finetuning

With the task $\mathcal{D}(f_*)$ and models ϕ_s, ϕ_w specified, we are ready to formalize the data and algorithms for FT.

We consider two sample sets drawn *i.i.d.* from $\mathcal{D}(f_*)$: a small labeled set $\tilde{\mathcal{S}} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i) | i \in [n]\} \sim \mathcal{D}(f_*)^n$ and a large sample set $\mathcal{S} = \{(\mathbf{x}_i, y_i) | i \in [N]\} \sim \mathcal{D}(f_*)^N$ where the labels y_i are inaccessible, denoting the unlabeled part as $\mathcal{S}_x = \{\mathbf{x}_i | i \in [N]\}$. The goal is to learn a function $f : \mathcal{X} \rightarrow \mathbb{R}$ using $\tilde{\mathcal{S}}$ and \mathcal{S}_x that generalizes well to $\mathcal{D}(f_*)$.

For $\tilde{\mathcal{S}}$, let $\tilde{\Phi}_w = [\phi_w(\tilde{\mathbf{x}}_1), \dots, \phi_w(\tilde{\mathbf{x}}_n)]^\top$, $\tilde{\Phi}_s = [\phi_s(\tilde{\mathbf{x}}_1), \dots, \phi_s(\tilde{\mathbf{x}}_n)]^\top \in \mathbb{R}^{n \times d}$ be the weak and strong features with associated labels $\tilde{\mathbf{y}} = [\tilde{y}_1, \dots, \tilde{y}_n]^\top \in \mathbb{R}^n$. Analogously for \mathcal{S} , let $\Phi_w = [\phi_w(\mathbf{x}_1), \dots, \phi_w(\mathbf{x}_N)]^\top$, $\Phi_s = [\phi_s(\mathbf{x}_1), \dots, \phi_s(\mathbf{x}_N)]^\top \in \mathbb{R}^{N \times d}$ be the weak and strong features with unknown labels $\mathbf{y} = [y_1, \dots, y_N]^\top \in \mathbb{R}^N$. For conciseness of notations, we introduce a mild regularity assumption on the ranks of these feature matrices.

Assumption 2.6 (Sufficient finetuning data). Assume $\tilde{\mathcal{S}}$ and \mathcal{S} are sufficiently large such that $\text{rank}(\tilde{\Phi}_w) = \text{rank}(\Phi_w) = d_w$ and $\text{rank}(\tilde{\Phi}_s) = \text{rank}(\Phi_s) = d_s$ almost surely².

Given regularization hyperparameters $\alpha_w, \alpha_{w2s}, \alpha_s, \alpha_c > 0$, we consider the following FT algorithms:

- (a) **Weak teacher model** $f_w(\mathbf{x}) = \phi_w(\mathbf{x})^\top \theta_w$ is supervisedly finetuned over $\tilde{\mathcal{S}}$:

$$\theta_w = \underset{\theta \in \mathbb{R}^d}{\text{argmin}} \frac{1}{n} \|\tilde{\Phi}_w \theta - \tilde{\mathbf{y}}\|_2^2 + \alpha_w \|\theta\|_2^2. \quad (2)$$

- (b) **W2S model** $f_{w2s}(\mathbf{x}) = \phi_s(\mathbf{x})^\top \theta_{w2s}$ is finetuned over the strong feature ϕ_s through \mathcal{S}_x and their pseudo-

² Assuming distributions of $\phi_w(\mathbf{x})$ and $\phi_s(\mathbf{x})$ are absolutely continuous with respect to the Lebesgue measure, for any $n \geq d_w$ and $n \geq d_s$, $\text{rank}(\tilde{\Phi}_w) = \text{rank}(\Phi_w) = d_w$ and $\text{rank}(\tilde{\Phi}_s) = \text{rank}(\Phi_s) = d_s$ almost surely (Vershynin, 2018, §3.3.1).

labels generated by the weak teacher model:

$$\theta_{w2s} = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \frac{1}{N} \|\Phi_s \theta - \Phi_w \theta_w\|_2^2 + \alpha_{w2s} \|\theta\|_2^2 \quad (3)$$

- (c) **Strong SFT model** $f_s(\mathbf{x}) = \phi_s(\mathbf{x})^\top \theta_s$ is a strong baseline where the strong feature ϕ_s is supervisedly finetuned over the small labeled set $\tilde{\mathcal{S}}$ directly:

$$\theta_s = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \frac{1}{n} \|\tilde{\Phi}_s \theta - \tilde{\mathbf{y}}\|_2^2 + \alpha_s \|\theta\|_2^2. \quad (4)$$

- (d) **Strong ceiling model** $f_c(\mathbf{x}) = \phi_s(\mathbf{x})^\top \theta_c$ is a reference for the ceiling performance where ϕ_s is supervisedly finetuned over $\mathcal{S} \cup \tilde{\mathcal{S}}$, assuming access to the unknown labels $\mathbf{y} = [y_1, \dots, y_N]^\top$:

$$\theta_c = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \frac{1}{n + N} \left\| \begin{bmatrix} \tilde{\Phi}_s \\ \Phi_s \end{bmatrix} \theta - \begin{bmatrix} \tilde{\mathbf{y}} \\ \mathbf{y} \end{bmatrix} \right\|_2^2 + \alpha_c \|\theta\|_2^2. \quad (5)$$

Remark 2.7 (Regularization prevents W2S from overfitting). As pointed out in (Burns et al., 2024), suitable regularization is crucial to prevent W2S from overfitting the weak teacher. For overparametrized problems³, even without explicit regularization, gradient descent implicitly biases toward the minimum ℓ_2 -norm solutions in the kernel regime (Woodworth et al., 2020), equivalent to solving Equations (2) to (5) with $\alpha_w, \alpha_{w2s}, \alpha_s, \alpha_c \rightarrow 0$. Therefore, we focus on ridgeless regression here under the idealized intrinsic dimension assumption in Definition 2.3. In Appendix C, we extend our analysis to the more general scenario: when Σ_s, Σ_w are not exactly low-rank, a careful choice of $\alpha_w, \alpha_{w2s} > 0$ brings a W2S generalization bound, Theorem C.4, that conveys the same message as Theorem 3.1 in the ridgeless case.

2.3. Metrics for W2S performance

In addition to the absolute generalization error of W2S, $\mathbf{ER}(f_{w2s})$, we quantify the W2S performance of f_{w2s} relative to f_w, f_s , and f_c through the following metrics:

- (a) **Performance gap recovery (PGR)** introduced in (Burns et al., 2024) measures the ratio between excess risk reductions from the weak teacher f_w of the W2S model f_{w2s} and the strong ceiling model f_c :

$$\mathbf{PGR} = \frac{\mathbf{ER}(f_w) - \mathbf{ER}(f_{w2s})}{\mathbf{ER}(f_w) - \mathbf{ER}(f_c)}. \quad (6)$$

³While the feature dimension d can be either larger (overparametrized) or smaller (underparametrized) than the sample sizes $n, N, n + N$, with the low intrinsic dimensions $d_s, d_w \ll d$, Equations (2) to (5) are always underdetermined.

In practice, $\mathbf{ER}(f_{w2s})$ typically falls between $\mathbf{ER}(f_c)$ and $\mathbf{ER}(f_w)$ (Burns et al., 2024). Therefore, it usually holds that $0 \leq \mathbf{PGR} \leq 1$. A higher \mathbf{PGR} indicates better W2S generalization: the W2S model f_{w2s} can recover more of the excess risk gap between the weak teacher f_w and the strong ceiling model f_c .

- (b) **Outperforming ratio (OPR)** compares excess risks of the strong baseline f_s and the W2S model f_{w2s} :

$$\mathbf{OPR} = \mathbf{ER}(f_s) / \mathbf{ER}(f_{w2s}). \quad (7)$$

A higher \mathbf{OPR} implies better W2S generalization: f_{w2s} outperforms f_s when $\mathbf{OPR} > 1$. This metric could be of interest in practice when the labeled samples $\tilde{\mathcal{S}}$ are limited – if $\mathbf{OPR} < 1$, SFT the strong model over $\tilde{\mathcal{S}}$ would be a better choice than W2S both in terms of generalization and computational efficiency.

3. Main results

In this section, we first analyze the generalization errors of W2S and its reference models in Section 3.1. Then in Section 3.2, we conduct a case study on the W2S performance in terms of the metrics introduced in Section 2.3.

3.1. Generalization errors

We start with the W2S model $f_{w2s}(\mathbf{x}) = \phi_s(\mathbf{x})^\top \theta_{w2s}$ finetuned as in (2), (3) with both $\alpha_w, \alpha_{w2s} \rightarrow 0$. For demonstration purposes, we consider an idealized Gaussian feature case in the main text, where the variance of f_{w2s} can be exactly characterized (instead of upper bounded)⁴.

Theorem 3.1 (W2S model (formally in B.1)). *Assuming Assumptions 2.1 and 2.6 and $\phi_w(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_d, \Sigma_w)$, for $n > d_w + 1$, $\mathbf{ER}(f_{w2s}) = \mathbf{Var}(f_{w2s}) + \mathbf{Bias}(f_{w2s})$ satisfies*

$$\mathbf{Var}(f_{w2s}) = \frac{\sigma^2}{n - d_w - 1} \left(d_{s \wedge w} + \frac{d_s}{N} (d_w - d_{s \wedge w}) \right),$$

$$\mathbf{Bias}(f_{w2s}) \leq \frac{\rho_w(n)}{n} + \frac{\rho_s(N)}{N} \leq \rho_w + \rho_s,$$

where the inequality for $\mathbf{Bias}(f_{w2s})$ is strict if $\rho_w(n)/n > 0$ and $d_s < d_w$.

Remark 3.2 (Discrepancy is virtue). Notice that $\mathbf{Var}(f_{w2s})$ consists of two terms. (a) In the overlapped subspace $\text{Range}(\Sigma_s) \cap \text{Range}(\Sigma_w)$ with correlation dimension $d_{s \wedge w}$, the variance $\sigma^2 d_{s \wedge w} / (n - d_w - 1)$ mimics that of the weak teacher, where more pseudo-labels N fail to reduce the variance. (b) Whereas variance in the subspace of discrepancy $\text{Range}(\Sigma_w) \setminus \text{Range}(\Sigma_s)$ takes the form $\sigma^2 (d_s / N) (d_w - d_{s \wedge w}) / (n - d_w - 1)$, reduced by a factor of d_s / N and vanishing as N grows.

⁴ The analogous generalization bound holds up to constants for sub-gaussian features in Assumption 2.1, see Theorem B.2.

As a reference, we also look into the weak teacher model $f_w(\mathbf{x}) = \phi_w(\mathbf{x})^\top \theta_w$ in (2) with $\alpha_w \rightarrow 0$:

Proposition 3.3 (Weak teacher (B.2)). *With Assumptions 2.1 and 2.6, $\text{ER}(f_w) = \text{Var}(f_w) + \text{Bias}(f_w)$ has*

$$\text{Var}(f_w) = \sigma^2 \frac{d_w}{n}, \quad \text{Bias}(f_w) = \frac{\rho_w(n)}{n} \leq \rho_w.$$

To measure the W2S performance in a relative sense, another two necessary references are the strong SFT baseline $f_s(\mathbf{x}) = \phi_s(\mathbf{x})^\top \theta_s$ in (4) and strong ceiling model $f_c(\mathbf{x}) = \phi_s(\mathbf{x})^\top \theta_c$ in (5), with both $\alpha_s, \alpha_c \rightarrow 0$:

Corollary 3.4 (Strong SFT and ceiling models). *Under Assumptions 2.1 and 2.6, $\text{ER}(f_s) = \text{Var}(f_s) + \text{Bias}(f_s)$ and $\text{ER}(f_c) = \text{Var}(f_c) + \text{Bias}(f_c)$ satisfy*

$$\text{Var}(f_s) = \sigma^2 \frac{d_s}{n}, \quad \text{Bias}(f_s) = \frac{\rho_s(n)}{n} \leq \rho_s,$$

$$\text{Var}(f_c) = \sigma^2 \frac{d_s}{N+n}, \quad \text{Bias}(f_c) = \frac{\rho_s(N+n)}{N+n} \leq \rho_s,$$

W2S in variance. Assuming $\rho_s + \rho_w \ll \sigma^2$ (Definition 2.2), variance dominates the generalization error. Theorem 3.1 and Proposition 3.3 suggest that W2S generalization occurs in variance, i.e. $\text{Var}(f_{w2s}) < \text{Var}(f_w)$, if

$$n(1 - d_s/N) > (d_w + 1)d_w/(d_w - d_{s \wedge w}),$$

which holds broadly when n, N are large enough. For example, with a low correlation dimension $d_{s \wedge w} < d_w$,

$$n > (d_w + 1)^2/(d_w - d_{s \wedge w}) \quad \text{and} \quad N \geq (d_w + 1)d_s$$

are sufficient. Notice that when $d_s < d_w$, $d_{s \wedge w} < d_w$ always holds because $d_{s \wedge w} \leq d_s < d_w$. Meanwhile, with a lower correlation dimension $d_{s \wedge w}$, W2S in variance is more pronounced and takes place at smaller n .

W2S in bias. Comparing the biases in Theorem 3.1 and Proposition 3.3, when the strong student has zero FT approximation error $\rho_s = 0$, and the FT approximation error of the weak teacher is non-negligible $\rho_w(n)/n > 0$, as long as the strong student has a lower intrinsic dimension $d_s < d_w$, W2S also enjoys a strictly lower bias than the weak teacher: $\text{Bias}(f_{w2s}) < \text{Bias}(f_w)$ ⁵.

3.2. W2S performance: a case study

With the generalization analysis, we are ready to take a closer look at the W2S performance in terms of **PGR** and **OPR** defined in Section 2.3.

⁵ Notice that quantifying such advantage of W2S in bias requires further assumptions on the downstream task $\mathcal{D}(f_*)$ and the covariance matrices Σ_w, Σ_s , analogous to the settings in (Ildiz et al., 2024; Wu & Sahai, 2024), which is deviating from our focus on variance but could be an interesting future direction.

Proposition 3.5 (**PGR** and **OPR** lower bounds (B.3)). *Given f_w, f_{w2s}, f_c , and f_s as in Theorem 3.1, Proposition 3.3, and Corollary 3.4, under Assumptions 2.1 and 2.6, assuming $\phi_w(\mathbf{x}) \sim \mathcal{N}(\theta_d, \Sigma_w)$ ⁴, with $n = d_w + q + 1$ for some constant $q \in \mathbb{N}$, we have **PGR** \geq*

$$1 - \frac{n}{q} \frac{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}{d_w} - \frac{n}{d_w} \frac{\rho_w + \rho_s}{\sigma^2},$$

and **OPR** \geq

$$\left(\frac{n}{q} \frac{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}{d_s} + \frac{n}{d_s} \frac{\rho_w + \rho_s}{\sigma^2} \right)^{-1}.$$

We recall from Section 2.3 that the larger **PGR** and **OPR** imply better W2S generalization. Then, a natural question hinted by Proposition 3.5 is *how do the sample sizes n, N affect the W2S performance?* The concrete answers to this question depend on the relative magnitude of the FT approximation errors and label noise, $(\rho_s + \rho_w)/\sigma^2$.

Case I: negligible FT approximation error. In the ideal case where the FT approximation errors are negligible compared to label noise, $(\rho_s + \rho_w)/\sigma^2 \rightarrow 0$, Proposition 3.5 suggests better lower bounds for **PGR**, **OPR** as n, N increase:

$$\text{PGR} \geq 1 - \frac{n}{n - d_w - 1} \frac{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}{d_w},$$

$$\text{OPR} \geq \frac{n - d_w - 1}{n} \frac{d_s}{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}.$$

Depending on $d_{s \wedge w}$, we have the following cases:

- When $d_{s \wedge w} > 0$, with sample sizes $n \gtrsim d_w$ and $N \gtrsim (d_w/d_{s \wedge w} - 1)d_s$, **PGR** $\geq 1 - O(d_s/N/d_w)$ and **OPR** $\geq \Omega(d_s/d_{s \wedge w})$ imply good W2S performance if $d_{s \wedge w} \ll \min\{d_s, d_w\}$.
- When $d_{s \wedge w} = 0$, a labeled sample size of $n \gtrsim d_w$ leads to **PGR** $\geq 1 - O(d_s/N)$ and **OPR** $\geq \Omega(N/d_w)$, implying good W2S performance when $N \gg \max\{d_w, d_s\}$.

Case II: small non-negligible FT approximation error.

In a more realistic scenario where $0 < (\rho_s + \rho_w)/\sigma^2 \ll 1$ is small but non-negligible, the trade-off between variance and bias brings about a non-monotonic scaling of the **PGR** and **OPR** lower bounds with respect to n :

Corollary 3.6 (Non-monotonic scaling w.r.t. n (B.4)). *For conciseness, denote $d_{w2s}(N) = d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N$ and $\varrho = (\rho_w + \rho_s)/\sigma^2$. Under the same setting as Proposition 3.5, with $n = d_w + q + 1$ for some $q \in \mathbb{N}$, the lower bounds of **PGR** and **OPR** are maximized when*

$q = \sqrt{(d_w + 1) d_{w2s}(N)/\rho}$, where

$$\text{PGR} \geq 1 - d_w^{-1} \left(\sqrt{d_{w2s}(N)} + \sqrt{\rho(d_w + 1)} \right)^2,$$

$$\text{OPR} \geq d_s \left(\sqrt{d_{w2s}(N)} + \sqrt{\rho(d_w + 1)} \right)^{-2}.$$

Such non-monotonic scaling for **PGR** with respect to n coincides with some empirical observations in (Burns et al., 2024) on NLP tasks. While the variance of f_{w2s} in Theorem 3.1 decreases monotonically as n grows, so do those of the reference models f_w , f_s , and f_c . With non-negligible FT approximation errors, as n increases, the **PGR** and **OPR** lower bounds decrease with the improvements in bias but increase with the improvements in variance. Therefore, the optimal n for the lower bounds of **PGR** and **OPR** is determined by the trade-off between variance and bias.

Assuming $\rho_s + \rho_w \ll \sigma^2$ in Definition 2.2, we have $\rho \ll 1$. Again, consider two cases depending on $d_{s \wedge w}$:

- If $d_{s \wedge w} > 0$, we have $d_{w2s}(N) \lesssim d_{s \wedge w}$ when $N \gtrsim (d_w/d_{s \wedge w} - 1)d_s$, implying large **PGR** and **OPR** when $d_{s \wedge w} \ll \min\{d_s, d_w\}$.
- If $d_{s \wedge w} = 0$, we have $d_{w2s}(N) = d_w d_s / N$, implying large **PGR** and **OPR** when $N \gg \max\{d_w, d_s\}$.

4. Experiments

We conduct experiments to validate the theoretical findings on both synthetic and real tasks. In this section, we focus on two illustrative settings: synthetic regression (Section 4.1) and real-world image regression (Section 4.2). For brevity, we defer more experiments on image and sentiment classification tasks to Appendices E.2 and E.3, respectively.

4.1. Synthetic regression

We start by grounding the theoretical framework introduced in Section 2 with synthetic regression tasks.

Setup. We concretize the downstream task $\mathcal{D}(f_*)$ as a regression problem over Gaussian features. Let $f_* : \mathbb{R}^d \rightarrow \mathbb{R}$ be a linear function in a high-dimensional feature space $d = 20,000$ of form $f_*(\mathbf{x}) = \mathbf{x}^\top \Lambda_*^{1/2} \boldsymbol{\theta}_*$ where $\Lambda_* = \text{diag}(\lambda_1^*, \dots, \lambda_d^*)$ is a diagonal matrix with a low rank $d_* = 300$ such that $\lambda_i^* = i^{-1}$ for $i \leq d_*$ and $\lambda_i^* = 0$ otherwise; and $\boldsymbol{\theta}_* \in \mathbb{R}^d$ is a random unit vector. Every sample $(\mathbf{x}, y) \sim \mathcal{D}(f_*)$ is generated by $\mathbf{x} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ and $y = f_*(\mathbf{x}) + z$ with $z \sim \mathcal{N}(0, \sigma^2)$. Given \mathbf{x} , the associated strong and weak features in Assumption 2.1 are generated by $\phi_s(\mathbf{x}) = \Sigma_s^{1/2} \mathbf{x}$ and $\phi_w(\mathbf{x}) = \Sigma_w^{1/2} \mathbf{x}$, with intrinsic dimensions $d_s = 100$ and $d_w = 200$ such that $\Sigma_s = \sum_{i=1}^{d_s} \lambda_i^* \mathbf{e}_i \mathbf{e}_i^\top$ and $\Sigma_w = \sum_{i=d_s-d_{s \wedge w}+1}^{d_w+d_{s \wedge w}} \lambda_i^* \mathbf{e}_i \mathbf{e}_i^\top$. For all synthetic experiments, we have $\rho_s + \rho_w < 0.0004$.

In the experiments, we vary $d_{s \wedge w}$ to control the student-

teacher correlation and σ^2 to control the dominance of variance over bias (characterized by ρ_s, ρ_w). Each error bar reflects the standard deviation over 40 runs.




Figure 1. Scaling for excess risks on the synthetic regression task in a variance-dominated regime with a low correlation dimension.




Figure 2. Scaling for excess risks on the synthetic regression task in a variance-dominated regime with a high correlation dimension.




Figure 3. Scaling for excess risks on the synthetic regression task when the variance is not dominant, $\sigma^2 \approx \rho_s + \rho_w$.

Scaling for generalization errors. Figures 1 to 3 show scaling for $\text{ER}(f_{w2s})$ (W2S), $\text{ER}(f_w)$ (Weak), $\text{ER}(f_s)$ (S-Baseline), and $\text{ER}(f_c)$ (S-Ceiling) with respect to the sample sizes n, N . The dashes show theoretical predictions in Theorem 3.1, Proposition 3.3, and Corollary 3.4, consistent with the empirical measurements shown in the solid lines. In particular, we consider three cases:

- Figure 1: When variance dominates ($\sigma^2 = 0.01 \gg \rho_w + \rho_s$), with a low correlation dimension $d_{s \wedge w} = 10$,

f_{w2s} outperforms both f_w and f_s for a moderate n and a large enough N . However, larger sample sizes do not necessarily lead to better W2S generalization in a relative sense. For example, when n keeps increasing, the strong baseline f_s eventually outperforms f_{w2s} .

- Figure 2: When variance dominates, with a high correlation dimension $d_{s \wedge w} = 90$, f_{w2s} still generalizes better than f_w but fails to outperform the strong baseline f_s .
- Figure 3: When the variance is low (not dominant, *e.g.* $\sigma^2 = 0.0004 \approx \rho_s + \rho_w$), f_{w2s} can fail to outperform f_w . This suggests that variance reduction is a key advantage of W2S over supervised FT.

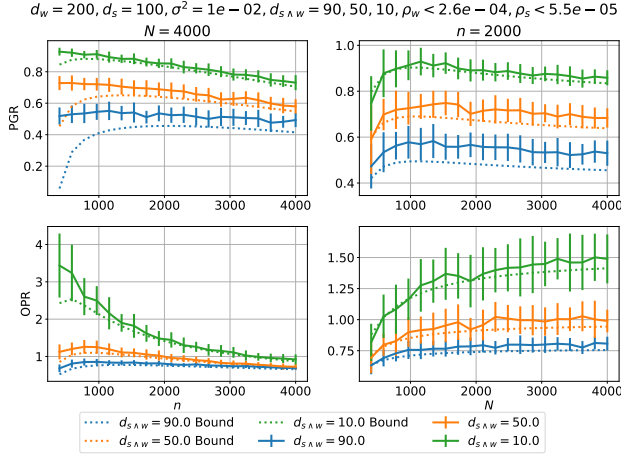


Figure 4. Scaling for PGR and OPR under different $d_{s \wedge w}$ on the synthetic regression task in a variance-dominated regime.

Scaling for PGR and OPR. Figure 4 show the scaling for PGR and OPR with respect to sample sizes n, N in the variance-dominated regime (with small non-negligible FT approximation errors), at three different correlation dimensions $d_{s \wedge w} = 90, 50, 10$. The solid and dashed lines represent the empirical measurements and lower bounds in (12), (13), respectively.

- Coinciding with the theoretical predictions in Corollary 3.6 and the performance gaps between W2S and the references in Figure 1, we observe that the relative W2S performance in terms of PGR and OPR can degenerate as n increases, while the larger N generally leads to better W2S generalization in the relative sense.
- The lower correlation dimension $d_{s \wedge w}$ leads to higher PGR and OPR, *i.e.* larger discrepancy between the strong and weak features improves W2S generalization.

4.2. UTKFace regression

Beyond the synthetic regression, we investigate W2S on a real-world image regression task – age estimation on the

UTKFace dataset (Zhang et al., 2017). Each error bar in this section reflects standard deviation of 10 runs.

Dataset. UTKFace (Aligned & Cropped) (Zhang et al., 2017) consists of 23,708 face images with age labels ranging from 0 to 116. We preprocess the images to 224×224 pixels and split the dataset into training and testing sets of sizes 20,000 and 3,708. Generalization errors are estimated with the mean squared error (MSE) over the test set.

Linear probing over pretrained features. We fix the strong student as CLIP ViT-B/32 (Radford et al., 2021) (CLIP-B32) and vary the weak teacher among the ResNet series (He et al., 2016) (ResNet18, ResNet34, ResNet50, ResNet101, ResNet152). We treat the backbones of these models (excluding the classification layers) as ϕ_s, ϕ_w and finetune them via linear probing. We use ridge regression with a small fixed regularization hyperparameter $\alpha_w, \alpha_{w2s}, \alpha_s, \alpha_c = 10^{-6}$, close to the machine epsilon of single precision floating point numbers.

Intrinsic dimension. The intrinsic dimensions d_w, d_s are measured based on the empirical covariance matrices Σ_w, Σ_s of the weak and strong features over the entire dataset (including training and testing). As mentioned in Footnote 1, these covariances generally have fast-decaying eigenvalues (but not exactly low-rank) in practice, effectively leading to low intrinsic dimensions under ridge regression. We estimate such low intrinsic dimensions as the minimum rank for the best low-rank approximation of Σ_w, Σ_s with a relative error in trace less than $\tau = 0.01$.

Correlation dimension. The pretrained feature dimensions (or the finetunable parameter counts) of the weak and strong models can be different in practice (see Appendix E.1, Table 1). We introduce an estimation for $d_{s \wedge w}$ in this case. Consider the (truncated) spectral decompositions $[\Sigma_s]_{d_s} = \mathbf{V}_s \Lambda_s \mathbf{V}_s^\top$ and $[\Sigma_w]_{d_w} = \mathbf{V}_w \Lambda_w \mathbf{V}_w^\top$ of two empirical covariances with different feature dimensions D_s, D_w such that $\mathbf{V}_s \in \mathbb{R}^{D_s \times d_s}$ and $\mathbf{V}_w \in \mathbb{R}^{D_w \times d_w}$ consists of the top d_s, d_w orthonormal eigenvectors, respectively. We estimate the correlation dimension $d_{s \wedge w}$ under different feature dimensions $D_s \neq D_w$ by matching the dimensions through a random unitary matrix (Vershynin, 2018) $\Gamma \in \mathbb{R}^{D_s \times D_w}$: $d_{s \wedge w} = \|\mathbf{V}_s^\top \Gamma \mathbf{V}_w\|_F^2$. This provides a good estimation for $d_{s \wedge w}$ because with low intrinsic dimensions $\max\{d_s, d_w\} \ll D_s, D_w$ in practice, mild dimension reduction through Γ well preserves the essential information in $\mathbf{V}_s, \mathbf{V}_w$.

Discrepancies lead to better W2S. Figure 5 shows the scaling of PGR and OPR with respect to the sample sizes n, N for different weak teachers in the ResNet series with respect to a fixed student, CLIP-B32. We first observe that




Figure 5. Scaling for **PGR** and **OPR** of different weak teachers with a fixed strong student on UTKFace. The legends show the comparison between $d_{s \wedge w}$ and d_w .

the relative W2S performance in terms of **PGR** and **OPR** is closely related to the correlation dimension $d_{s \wedge w}$ and the intrinsic dimensions d_s, d_w .

- When the strong student has a lower intrinsic dimension than the weak teacher (as widely observed in practice (Aghajanyan et al., 2021)), *i.e.* $d_s < d_w$, the relative W2S performance tends to be better than when $d_s > d_w$.
- The relative W2S performance tends to be better when $d_{s \wedge w}/d_w$ is lower, *i.e.* the larger discrepancy between weak and strong features leads to better W2S generalization.

Meanwhile, both **PGR** and **OPR** scale inversely with the labeled sample size n and exhibit diminishing return with respect to the increasing pseudolabel size N , consistent with the theoretical predictions in Corollary 3.6 and the synthetic experiments in Figure 4.

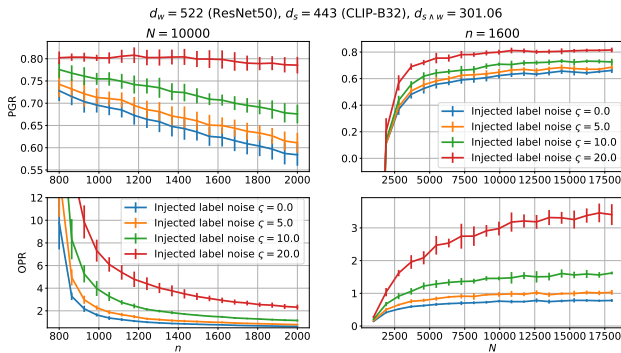


Figure 6. Scaling for **PGR** and **OPR** on UTKFace with injected label noise: $y_i \leftarrow y_i + \zeta_i$ where $\zeta_i \sim \mathcal{N}(0, \zeta^2)$ *i.i.d.*

Variance reduction is a key advantage of W2S. To investigate the impact of variance on W2S generalization, we inject noise to the training label by $y_i \leftarrow y_i + \zeta_i$ where $\zeta_i \sim \mathcal{N}(0, \zeta^2)$ *i.i.d.*, and ζ controls the injected labels noise

level. In Figure 6, we show the scaling for **PGR** and **OPR** with respect to the sample sizes n, N under different noise levels ζ . We observe that the relative W2S performance in terms of **PGR** and **OPR** improves as the noise level ζ increases. This provides empirical evidence that variance reduction is a key advantage of W2S over supervised FT, highlighting the importance of understanding the mechanisms of W2S in the variance-dominated regime.

5. Limitations and future directions

In this work, we introduce a theoretical framework for understanding the mechanism of weak-to-strong (W2S) generalization in the variance-dominated regime where both the student and teacher have sufficient capacities for the downstream task. Leveraging the low intrinsic dimensionality of finetuning (FT), we characterize model capacities from three perspectives: FT approximation errors for “accuracy”, intrinsic dimensions for “complexity”, and student-teacher correlation for “alignment”. Our analysis shows that W2S generalization is driven by variance reduction in the discrepancy between the weak teacher and strong student features. This generalization analysis is followed by a case study on the relative W2S performance in terms of performance gap recovery (PGR) and outperforming ratio (OPR). We show that while larger sample sizes imply better W2S generalization in an absolute sense, the relative W2S performance can degenerate as the sample size increases. Our results provide theoretical insights into the choice of weak teachers and sample sizes in W2S pipelines.

An interesting implication of our analysis is that the mechanism of W2S may differ as the balance between variance and bias shifts. In the variance-dominated regime studied in this work, W2S can benefit from a lower intrinsic dimension of the strong student due to the resulting variance reduction in the subspace of discrepancy from the weak teacher. In contrast, in the bias-dominated regime, the lower approximation error of the strong student is generally brought by the larger “capacity” of the strong model corresponding to a higher intrinsic dimension (Ildiz et al., 2024; Wu & Sahai, 2024). This calls for future studies on unified views and transitions between the two regimes, which will provide a more comprehensive understanding of W2S. Toward this goal, a limitation of our analysis is the quantification of the advantage of W2S in bias (see Footnote 5), which could be a promising next step.

Acknowledgements

The authors would like to thank Denny Wu for insightful discussions. The experiments are supported by the PLI computing cluster. YD acknowledges support of NYU Courant Instructorship. JDL acknowledges support of Open Philanthropy, NSF IIS 2107304, NSF CCF 2212262, NSF CA-

REER Award 2144994, and NSF CCF 2019844. This material is based upon work supported by the U.S. Department of Energy, Office of Science Energy Earthshot Initiative as part of the project “Learning reduced models under extreme data conditions for design and rapid decision-making in complex systems” under Award #DE-SC0024721.

Impact Statement

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- Aghajanyan, A., Gupta, S., and Zettlemoyer, L. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 7319–7328, 2021.
- Amari, S.-i., Ba, J., Grosse, R., Li, X., Nitanda, A., Suzuki, T., Wu, D., and Xu, J. When does preconditioning help or hurt generalization? *arXiv preprint arXiv:2006.10732*, 2020.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Bahri, Y., Dyer, E., Kaplan, J., Lee, J., and Sharma, U. Explaining neural scaling laws. *Proceedings of the National Academy of Sciences*, 121(27):e2311878121, 2024.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828, 2013.
- Björck, A. and Golub, G. H. Numerical methods for computing angles between linear subspaces. *Mathematics of computation*, 27(123):579–594, 1973.
- Blum, A. and Mitchell, T. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 92–100, 1998.
- Borup, K. and Andersen, L. N. Self-distillation for gaussian process regression and classification. *arXiv preprint arXiv:2304.02641*, 2023.
- Burns, C., Izmailov, P., Kirchner, J. H., Baker, B., Gao, L., Aschenbrenner, L., Chen, Y., Ecoffet, A., Joglekar, M., Leike, J., et al. Weak-to-strong generalization: eliciting strong capabilities with weak supervision. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 4971–5012, 2024.
- Cai, T., Gao, R., Lee, J., and Lei, Q. A theory of label propagation for subpopulation shift. In *International Conference on Machine Learning*, pp. 1170–1182. PMLR, 2021.
- Charikar, M., Pabbaraju, C., and Shiragur, K. Quantifying the gain in weak-to-strong generalization. *Advances in neural information processing systems*, 37:126474–126499, 2024.
- Clark, K., Luong, M.-T., Le, Q. V., and Manning, C. D. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.
- Das, R. and Sanghavi, S. Understanding self-distillation in the presence of label noise. In *International Conference on Machine Learning*, pp. 7102–7140. PMLR, 2023.
- Dong, Y. and Martinsson, P.-G. Simpler is better: a comparative study of randomized pivoting algorithms for cur and interpolative decompositions. *Advances in Computational Mathematics*, 49(4):66, 2023.
- Dong, Y., Xie, Y., and Ward, R. Adaptively weighted data augmentation consistency regularization for robust optimization under concept shift. In *International Conference on Machine Learning*, pp. 8296–8316. PMLR, 2023.
- Dong, Y., Martinsson, P.-G., and Nakatsukasa, Y. Efficient bounds and estimates for canonical angles in randomized subspace approximations. *SIAM Journal on Matrix Analysis and Applications*, 45(4):1978–2006, 2024a.
- Dong, Y., Miller, K., Lei, Q., and Ward, R. Cluster-aware semi-supervised learning: relational knowledge distillation provably learns clustering. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Dong, Y., Phan, H., Pan, X., and Lei, Q. Sketchy moment matching: Toward fast and provable data selection for finetuning. *arXiv preprint arXiv:2407.06120*, 2024c.
- Frei, S., Zou, D., Chen, Z., and Gu, Q. Self-training converts weak learners to strong learners in mixture models. In *International Conference on Artificial Intelligence and Statistics*, pp. 8003–8021. PMLR, 2022.
- Furlanello, T., Lipton, Z., Tschannen, M., Itti, L., and Anandkumar, A. Born again neural networks. In *International*

- tional conference on machine learning*, pp. 1607–1616. PMLR, 2018.
- Goel, S., Struber, J., Auzina, I. A., Chandra, K. K., Kumaraguru, P., Kiela, D., Prabhu, A., Bethge, M., and Geiping, J. Great models think alike and this undermines ai oversight. *arXiv preprint arXiv:2502.04313*, 2025.
- Golub, G. H. and Van Loan, C. F. *Matrix computations*. JHU press, 2013.
- Gou, J., Yu, B., Maybank, S. J., and Tao, D. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- Guo, J., Chen, H., Wang, C., Han, K., Xu, C., and Wang, Y. Vision superalignment: Weak-to-strong generalization for vision foundation models. *arXiv preprint arXiv:2402.03749*, 2024.
- Guo, Y. and Yang, Y. Improving weak-to-strong generalization with reliability-aware alignment. *arXiv preprint arXiv:2406.19032*, 2024.
- Halko, N., Martinsson, P.-G., and Tropp, J. A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- He, T., Zhang, Z., Zhang, H., Zhang, Z., Xie, J., and Li, M. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 558–567, 2019.
- Hinton, G. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Huang, W., Yi, M., Zhao, X., and Jiang, Z. Towards the generalization of contrastive self-supervised learning. *arXiv preprint arXiv:2111.00743*, 2021.
- Ildiz, M. E., Gozeten, H. A., Taga, E. O., Mondelli, M., and Oymak, S. High-dimensional analysis of knowledge distillation: Weak-to-strong generalization and scaling laws. *arXiv preprint arXiv:2410.18837*, 2024.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- Johnson, W. B. Extensions of lipshitz mapping into hilbert space. In *Conference modern analysis and probability, 1984*, pp. 189–206, 1984.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lang, H., Sontag, D., and Vijayaraghavan, A. Theoretical analysis of weak-to-strong generalization. *arXiv preprint arXiv:2405.16043*, 2024.
- Lee, D.-H. et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, pp. 896. Atlanta, 2013.
- Li, C., Farkhoor, H., Liu, R., and Yosinski, J. Measuring the intrinsic dimension of objective landscapes. In *International Conference on Learning Representations*, 2018.
- Li, M., Zhang, Y., He, S., Li, Z., Zhao, H., Wang, J., Cheng, N., and Zhou, T. Superfiltering: Weak-to-strong data filtering for fast instruction-tuning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14255–14273, 2024.
- Liu, Y. and Alahi, A. Co-supervised learning: Improving weak-to-strong generalization with hierarchical mixture of experts. *arXiv preprint arXiv:2402.15505*, 2024.
- Malladi, S., Wetteg, A., Yu, D., Chen, D., and Arora, S. A kernel-based view of language model fine-tuning. In *International Conference on Machine Learning*, pp. 23610–23641. PMLR, 2023.
- Mardia, K. V., Kent, J. T., and Taylor, C. C. *Multivariate analysis*, volume 88. John Wiley & Sons, 2024.
- Medvedev, M., Lyu, K., Yu, D., Arora, S., Li, Z., and Srebro, N. Weak-to-strong generalization even in random feature networks, provably. *arXiv preprint arXiv:2503.02877*, 2025.
- Mobahi, H., Farajtabar, M., and Bartlett, P. Self-distillation amplifies regularization in hilbert space. *Advances in Neural Information Processing Systems*, 33:3351–3361, 2020.
- Muthukumar, V., Narang, A., Subramanian, V., Belkin, M., Hsu, D., and Sahai, A. Classification vs regression in overparameterized regimes: Does the loss function matter? *Journal of Machine Learning Research*, 22(222): 1–69, 2021.
- Nagarajan, V., Menon, A. K., Bhojanapalli, S., Mobahi, H., and Kumar, S. On student-teacher deviations in distillation: does it pay to disobey? *Advances in Neural Information Processing Systems*, 36:5961–6000, 2023.
- Ojha, U., Li, Y., Sundara Rajan, A., Liang, Y., and Lee, Y. J. What knowledge gets distilled in knowledge distillation?

- Advances in Neural Information Processing Systems*, 36: 11037–11048, 2023.
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research Journal*, pp. 1–31, 2024.
- Oymak, S. and Gulcu, T. C. A theoretical characterization of semi-supervised learning with self-training for gaussian mixture models. In *International Conference on Artificial Intelligence and Statistics*, pp. 3601–3609. PMLR, 2021.
- Pang, B. and Lee, L. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pp. 115–124, 2005.
- Pareek, D., Du, S. S., and Oh, S. Understanding the gains from repeated self-distillation. *arXiv preprint arXiv:2407.04600*, 2024.
- Phuong, M. and Lampert, C. Towards understanding knowledge distillation. In *International conference on machine learning*, pp. 5142–5151. PMLR, 2019.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Scudder, H. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965.
- Shen, K., Jones, R. M., Kumar, A., Xie, S. M., HaoChen, J. Z., Ma, T., and Liang, P. Connect, not collapse: Explaining contrastive learning for unsupervised domain adaptation. In *International conference on machine learning*, pp. 19847–19878. PMLR, 2022.
- Shin, C., Cooper, J., and Sala, F. Weak-to-strong generalization through the data-centric lens. *arXiv preprint arXiv:2412.03881*, 2024.
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A. Y., and Potts, C. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pp. 1631–1642, 2013.
- Somerstep, S., Polo, F. M., Banerjee, M., Ritov, Y., Yurochkin, M., and Sun, Y. A statistical framework for weak-to-strong generalization. *arXiv preprint arXiv:2405.16236*, 2024.
- Spigler, S., Geiger, M., and Wyart, M. Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124001, 2020.
- Stanton, S., Izmailov, P., Kirichenko, P., Alemi, A. A., and Wilson, A. G. Does knowledge distillation really work? *Advances in Neural Information Processing Systems*, 34: 6906–6919, 2021.
- Sun, Z., Yu, L., Shen, Y., Liu, W., Yang, Y., Welleck, S., and Gan, C. Easy-to-hard generalization: Scalable alignment beyond human supervision. In *Annual Conference on Neural Information Processing Systems*, 2024.
- Tenenbaum, J. B., Silva, V. d., and Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- Turc, I., Chang, M.-W., Lee, K., and Toutanova, K. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*, 2019.
- Udell, M. and Townsend, A. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science*, 1(1):144–160, 2019.
- Van der Maaten, L. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Vershynin, R. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Wei, A., Hu, W., and Steinhardt, J. More than a toy: Random matrix models predict how real-world neural representations generalize. In *International Conference on Machine Learning*, pp. 23549–23588. PMLR, 2022.
- Wei, C., Shen, K., Chen, Y., and Ma, T. Theoretical analysis of self-training with deep networks on unlabeled data. *arXiv preprint arXiv:2010.03622*, 2020.
- Wishart, J. The generalised product moment distribution in samples from a normal multivariate population. *Biometrika*, pp. 32–52, 1928.
- Woodruff, D. P. et al. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- Woodworth, B., Gunasekar, S., Lee, J. D., Moroshko, E., Savarese, P., Golan, I., Soudry, D., and Srebro, N. Kernel and rich regimes in overparametrized models. In

- Conference on Learning Theory*, pp. 3635–3673. PMLR, 2020.
- Wu, D. X. and Sahai, A. Provable weak-to-strong generalization via benign overfitting. *arXiv preprint arXiv:2410.04638*, 2024.
- Xia, M., Malladi, S., Gururangan, S., Arora, S., and Chen, D. Less: selecting influential data for targeted instruction tuning. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 54104–54132, 2024.
- Yang, S., Dong, Y., Ward, R., Dhillon, I. S., Sanghavi, S., and Lei, Q. Sample efficiency of data augmentation consistency regularization. In *International Conference on Artificial Intelligence and Statistics*, pp. 3825–3853. PMLR, 2023.
- Yang, W., Shen, S., Shen, G., Yao, W., Liu, Y., Gong, Z., Lin, Y., and Wen, J.-R. Super (ficial)-alignment: Strong models may deceive weak models in weak-to-strong generalization. *arXiv preprint arXiv:2406.11431*, 2024a.
- Yang, Y., Ma, Y., and Liu, P. Weak-to-strong reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 8350–8367, 2024b.
- Yang, Z., Yu, Y., You, C., Steinhardt, J., and Ma, Y. Rethinking bias-variance trade-off for generalization of neural networks. In *International Conference on Machine Learning*, pp. 10767–10777. PMLR, 2020.
- Yao, W., Yang, W., Wang, Z., Lin, Y., and Liu, Y. Understanding the capabilities and limitations of weak-to-strong generalization. *arXiv preprint arXiv:2502.01458*, 2025.
- Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., and Ma, K. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 3713–3722, 2019.
- Zhang, L., Bao, C., and Ma, K. Self-distillation: Towards efficient and compact neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4388–4403, 2021.
- Zhang, Z., Song, Y., and Qi, H. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5810–5818, 2017.

Appendix

Contents

A Additional related works	14
B Proofs in Section 3	14
B.1 Proof of Theorem 3.1	15
B.2 Proof of Proposition 3.3 and Corollary 3.4	20
B.3 Proof of Proposition 3.5	21
B.4 Proof of Corollary 3.6	22
C Ridge regression analysis	22
D Canonical angles	28
E Additional experiments	29
E.1 Additional experiments and details on UTKFace regression	29
E.2 Experiments on image classification	31
E.3 Experiments on sentiment classification	32

A. Additional related works

Knowledge distillation. Knowledge distillation (KD) (Hinton, 2015; Gou et al., 2021) is closely connected to W2S generalization regarding the teacher-student setup, while W2S reverts the capacities of teacher and student in KD. In KD, a strong teacher model guides a weak student model to learn the teacher’s knowledge. In contrast, W2S generalization occurs when a strong student model surpasses a weak teacher model under weak supervision. Phuong & Lampert (2019); Stanton et al. (2021); Ojha et al. (2023); Nagarajan et al. (2023); Dong et al. (2024b); Ildiz et al. (2024) conducted rigorous statistical analyses for the student’s generalization from knowledge distillation. From the analysis perspective, a key difference between KD and W2S is that W2S is usually analyzed in the context of finetuning since the notions of “weak” and “strong” are built upon pretraining. This finetuning perspective introduces distinct angles from KD for examining intrinsic dimension (Li et al., 2018) and student-teacher correlation in W2S.

Self-distillation and self-training. In contrast to W2S, which considers distinct student and teacher models, self-distillation (Zhang et al., 2019; 2021) and related paradigms such as Born-Again Networks (Furlanello et al., 2018) use the same or progressively refined architectures to iteratively distill knowledge from a “previous version” of the model. There have been extensive theoretical analyses toward understanding the mechanism behind self-distillation (Mobahi et al., 2020; Das & Sanghavi, 2023; Borup & Andersen, 2023; Pareek et al., 2024).

Self-training (Scudder, 1965; Lee et al., 2013) is a closely related method to self-distillation that takes a single model’s confident predictions to create pseudo-labels for unlabeled data and refines that model iteratively. Wei et al. (2020); Oymak & Gulcu (2021); Frei et al. (2022) provide theoretical insights into the generalization of self-training. In particular, Wei et al. (2020) introduced a theoretical framework based on neighborhood expansion, which was later on extended to various settings of weakly supervised learning, including domain adaptation (Cai et al., 2021), contrastive learning (Shen et al., 2022; Huang et al., 2021), consistency regularization (Yang et al., 2023; Dong et al., 2023), and recently weak-to-strong generalization (Lang et al., 2024; Shin et al., 2024).

B. Proofs in Section 3

Lemma B.1. *Given the FT approximation errors ρ_s and ρ_w in Definition 2.2, we have*

$$\rho_s(n) \leq n\rho_s \quad \text{and} \quad \rho_w(n) \leq n\rho_w \quad \forall n \in \mathbb{N}.$$

Proof of Lemma B.1. Let $\boldsymbol{\theta}_* = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_w(\mathbf{x})^\top \boldsymbol{\theta} - f_*(\mathbf{x}))^2]$ such that

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_w(\mathbf{x})^\top \boldsymbol{\theta}_* - f_*(\mathbf{x}))^2] = \rho_w.$$

Then, by observing that conditioned on \mathbf{X} ,

$$\phi_w(\mathbf{X})^\dagger f_*(\mathbf{X}) = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\phi_w(\mathbf{X})\boldsymbol{\theta} - f_*(\mathbf{X})\|_2^2,$$

we have

$$\begin{aligned} \rho_w(n) &= \mathbb{E}_{\mathbf{X} \sim \mathcal{D}^n} [\|\phi_w(\mathbf{X})\phi_w(\mathbf{X})^\dagger f_*(\mathbf{X}) - f_*(\mathbf{X})\|_2^2] \\ &\leq \mathbb{E}_{\mathbf{X} \sim \mathcal{D}^n} [\|\phi_w(\mathbf{X})\boldsymbol{\theta}_* - f_*(\mathbf{X})\|_2^2] \\ &= n \mathbb{E}_{\mathbf{X} \sim \mathcal{D}^n} \left[\frac{1}{n} \|\phi_w(\mathbf{X})\boldsymbol{\theta}_* - f_*(\mathbf{X})\|_2^2 \right] \\ &= n \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_w(\mathbf{x})^\top \boldsymbol{\theta}_* - f_*(\mathbf{x}))^2] \\ &= n \rho_w. \end{aligned}$$

The proof for $\rho_s(n)$ follows analogously. \square

B.1. Proof of Theorem 3.1

Theorem B.2 (Formal restatement of Theorem 3.1). *Consider $f_{w2s}(\mathbf{x}) = \phi_s(\mathbf{x})^\top \boldsymbol{\theta}_{w2s}$ finetuned as in (2), (3) with both $\alpha_w, \alpha_{w2s} \rightarrow 0$. Under Assumptions 2.1 and 2.6, when $n \geq \Omega(d_w)$, the excess risk $\mathbf{ER}(f_{w2s}) = \mathbf{Var}(f_{w2s}) + \mathbf{Bias}(f_{w2s})$ satisfies*

$$\begin{aligned} \mathbf{Bias}(f_{w2s}) &\leq \frac{\rho_w(n)}{n} + \frac{\rho_s(N)}{N} \leq \rho_w + \rho_s, \\ \mathbf{Var}(f_{w2s}) &\lesssim \frac{\sigma^2}{n} \left(d_{s \wedge w} + \frac{d_s}{N} (d_w - d_{s \wedge w}) \right). \end{aligned}$$

In particular, when $\rho_w(n)/n > 0$ and $d_s < d_w$, the inequality for $\mathbf{Bias}(f_{w2s})$ is strict.

Moreover, when $\phi_w(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_d, \boldsymbol{\Sigma}_w)$, for any $n > d_w + 1$, we have

$$\mathbf{Var}(f_{w2s}) = \frac{\sigma^2}{n - d_w - 1} \left(d_{s \wedge w} + \frac{d_s}{N} (d_w - d_{s \wedge w}) \right).$$

Proof of Theorem 3.1 and Theorem B.2. We first observe that the solution of (2) as $\alpha_w \rightarrow 0$ is given by

$$\boldsymbol{\theta}_w = \tilde{\boldsymbol{\Phi}}_w^\dagger \tilde{\mathbf{y}} = \tilde{\boldsymbol{\Phi}}_w^\dagger (\tilde{\mathbf{f}}_* + \tilde{\mathbf{z}}),$$

where $\tilde{\mathbf{z}} \sim \mathcal{N}(\mathbf{0}_n, \sigma^2 \mathbf{I}_n)$. Meanwhile, the solution of (3) as $\alpha_{w2s} \rightarrow 0$ is given by

$$\boldsymbol{\theta}_{w2s} = \boldsymbol{\Phi}_s^\dagger \boldsymbol{\Phi}_w \boldsymbol{\theta}_w = \boldsymbol{\Phi}_s^\dagger \boldsymbol{\Phi}_w \tilde{\boldsymbol{\Phi}}_w^\dagger (\tilde{\mathbf{f}}_* + \tilde{\mathbf{z}}).$$

Then, the excess risk of f_{w2s} can be decomposed into variance and bias as follows:

$$\begin{aligned} \mathbf{ER}(f_{w2s}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\mathbb{E}_{f_{w2s}} [(f_{w2s}(\mathbf{x}) - f_*(\mathbf{x}))^2]] \\ &= \mathbb{E}_{\mathcal{S}_x} \left[\mathbb{E}_{\tilde{\mathcal{S}}} \left[\frac{1}{N} \|\boldsymbol{\Phi}_s \boldsymbol{\theta}_{w2s} - \mathbf{f}_*\|_2^2 \right] \right] \\ &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| (\boldsymbol{\Phi}_s \boldsymbol{\Phi}_s^\dagger \boldsymbol{\Phi}_w \tilde{\boldsymbol{\Phi}}_w^\dagger \tilde{\mathbf{f}}_* - \mathbf{f}_*) + \boldsymbol{\Phi}_s \boldsymbol{\Phi}_s^\dagger \boldsymbol{\Phi}_w \tilde{\boldsymbol{\Phi}}_w^\dagger \tilde{\mathbf{z}} \right\|_2^2 \right] \\ &= \underbrace{\frac{1}{N} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \boldsymbol{\Phi}_s \boldsymbol{\Phi}_s^\dagger \boldsymbol{\Phi}_w \tilde{\boldsymbol{\Phi}}_w^\dagger \tilde{\mathbf{z}} \right\|_2^2 \right]}_{\mathbf{Var}(f_{w2s})} + \underbrace{\frac{1}{N} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \boldsymbol{\Phi}_s \boldsymbol{\Phi}_s^\dagger \boldsymbol{\Phi}_w \tilde{\boldsymbol{\Phi}}_w^\dagger \tilde{\mathbf{f}}_* - \mathbf{f}_* \right\|_2^2 \right]}_{\mathbf{Bias}(f_{w2s})}. \end{aligned}$$

Bias. For the bias term, by observing that $\mathbf{P}_s = \Phi_s \Phi_s^\dagger$ is an $N \times N$ orthogonal projection, we can decompose the bias term as

$$\text{Bias}(f_{w2s}) = \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{P}_s \left(\Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \mathbf{f}_* \right) \right\|_2^2 \right] + \frac{1}{N} \mathbb{E}_{\mathcal{S}_x} \left[\left\| (\mathbf{I}_N - \mathbf{P}_s) \mathbf{f}_* \right\|_2^2 \right],$$

where $\mathbb{E}_{\mathcal{S}_x} \left[\left\| (\mathbf{I}_N - \mathbf{P}_s) \mathbf{f}_* \right\|_2^2 \right] = \rho_s(N)$ by Definition 2.2.

For the first term, we observe that

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{P}_s \left(\Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \mathbf{f}_* \right) \right\|_2^2 \right] &\leq \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \mathbf{f}_* \right\|_2^2 \right] \\ &= \mathbb{E}_{\tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_*} \left[\mathbb{E}_{(\Phi_w, \mathbf{f}_*)} \left[\frac{1}{N} \left\| \Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \mathbf{f}_* \right\|_2^2 \mid \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* \right] \right] \\ &= \mathbb{E}_{\tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_*} \left[\mathbb{E}_{(\tilde{\Phi}_w, \tilde{\mathbf{f}}_*)} \left[\frac{1}{n} \left\| \tilde{\Phi}_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \tilde{\mathbf{f}}_* \right\|_2^2 \mid \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* \right] \right] \\ &\leq \mathbb{E}_{\tilde{\mathcal{S}}} \left[\frac{1}{n} \left\| \tilde{\Phi}_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \tilde{\mathbf{f}}_* \right\|_2^2 \right] = \frac{\rho_w(n)}{n}. \end{aligned}$$

Notice that when $\rho_w(n)/n > 0$, this inequality is strict if $d_s < d_w$, where $\Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{f}}_* - \tilde{\mathbf{f}}_* \notin \text{Range}(\Phi_s)$ almost surely.

Overall, we have

$$\text{Bias}(f_{w2s}) \leq \frac{\rho_w(n)}{n} + \frac{\rho_s(N)}{N} \leq \rho_w + \rho_s,$$

where the second inequality follows from Lemma B.1.

Variance. For the variance term, we observe that

$$\begin{aligned} \text{Var}(f_{w2s}) &= \frac{1}{N} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \mathbf{P}_s \Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{z}} \right\|_2^2 \right] \\ &= \frac{1}{N} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\text{tr} \left(\Phi_w^\top \mathbf{P}_s \Phi_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{z}} \tilde{\mathbf{z}}^\top (\tilde{\Phi}_w^\dagger)^\top \right) \right] \\ &= \frac{\sigma^2}{N} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\text{tr} \left(\Phi_w^\top \mathbf{P}_s \Phi_w (\tilde{\Phi}_w^\top \tilde{\Phi}_w)^\dagger \right) \right], \end{aligned}$$

which implies

$$\text{Var}(f_{w2s}) = \frac{\sigma^2}{N} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\Sigma_w^{-1/2} \Phi_w^\top \mathbf{P}_s \Phi_w \Sigma_w^{-1/2} \right] \mathbb{E}_{\tilde{\mathcal{S}}} \left[\left(\Sigma_w^{-1/2} \tilde{\Phi}_w^\top \tilde{\Phi}_w \Sigma_w^{-1/2} \right)^\dagger \right] \right). \quad (8)$$

Recall the spectral decomposition $\Sigma_w = \mathbf{V}_w \Lambda_w \mathbf{V}_w^\top$. Since $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\phi_w(\mathbf{x}) \phi_w(\mathbf{x})^\top] = \Sigma_w$, for each $\mathbf{x} \sim \mathcal{D}$, we can write $\phi_w(\mathbf{x}) = \Sigma_w^{1/2} \gamma$, where $\gamma \in \mathbb{R}^d$ is an independent random vector that is zero-mean and isotropic (*i.e.* $\mathbb{E}[\gamma] = \mathbf{0}_d$ and $\mathbb{E}[\gamma \gamma^\top] = \mathbf{I}_d$). The same holds for $\Sigma_s = \mathbf{V}_s \Lambda_s \mathbf{V}_s^\top$ and $\phi_s(\mathbf{x}) = \Sigma_s^{1/2} \gamma$.

Then, for \mathcal{S} and $\tilde{\mathcal{S}}$, there exist independent random matrices $\Gamma = [\gamma_1, \dots, \gamma_N]^\top \in \mathbb{R}^{N \times d}$ and $\tilde{\Gamma} = [\tilde{\gamma}_1, \dots, \tilde{\gamma}_n]^\top \in \mathbb{R}^{n \times d}$ consisting of *i.i.d.* zero-mean isotropic rows such that

$$\begin{aligned} \Phi_w \Sigma_w^{-1/2} &= \Gamma \Sigma_w^{1/2} \Sigma_w^{-1/2} = \Gamma \mathbf{V}_w \mathbf{V}_w^\top, \\ \tilde{\Phi}_w \Sigma_w^{-1/2} &= \tilde{\Gamma} \Sigma_w^{1/2} \Sigma_w^{-1/2} = \tilde{\Gamma} \mathbf{V}_w \mathbf{V}_w^\top, \\ \Phi_s \Sigma_s^{-1/2} &= \Gamma \Sigma_s^{1/2} \Sigma_s^{-1/2} = \Gamma \mathbf{V}_s \mathbf{V}_s^\top, \\ \tilde{\Phi}_s \Sigma_s^{-1/2} &= \tilde{\Gamma} \Sigma_s^{1/2} \Sigma_s^{-1/2} = \tilde{\Gamma} \mathbf{V}_s \mathbf{V}_s^\top. \end{aligned} \quad (9)$$

Let $\mathbf{\Gamma}_w = \mathbf{\Gamma} \mathbf{V}_w \in \mathbb{R}^{N \times d_w}$ and $\tilde{\mathbf{\Gamma}}_w = \tilde{\mathbf{\Gamma}} \mathbf{V}_w \in \mathbb{R}^{n \times d_w}$. We observe that

$$\mathbb{E}_{\tilde{\mathcal{S}}} \left[\left(\mathbf{\Sigma}_w^{-1/2} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w \mathbf{\Sigma}_w^{-1/2} \right)^\dagger \right] = \mathbb{E}_{\tilde{\mathcal{S}}} \left[\left(\mathbf{V}_w \tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w \mathbf{V}_w^\top \right)^\dagger \right] = \mathbf{V}_w \mathbb{E}_{\tilde{\mathcal{S}}} \left[\left(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w \right)^\dagger \right] \mathbf{V}_w^\top.$$

Now, we consider the following two cases for the feature distribution of $\phi_w(\mathbf{x})$, corresponding to the distribution of $\mathbf{\Gamma}_w$ and $\tilde{\mathbf{\Gamma}}_w$:

- (a) **Gaussian features:** In Theorem 3.1, assuming $\phi_w(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_d, \mathbf{\Sigma}_w)$ such that $\tilde{\mathbf{\Gamma}}_w$ consists of *i.i.d.* Gaussian rows, we have $\tilde{\gamma}_i \sim \mathcal{N}(\mathbf{0}_{d_w}, \mathbf{I}_{d_w})$. Notice that under the assumption $n > d_w + 1$, $\text{rank}(\tilde{\mathbf{\Gamma}}_w) = d_w$ almost surely, and therefore $\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w$ is invertible.

Meanwhile, with $\tilde{\gamma}_i \sim \mathcal{N}(\mathbf{0}_{d_w}, \mathbf{I}_{d_w})$ for all $i \in [n]$, $(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w) \sim \mathcal{W}(\mathbf{I}_{d_w}, n)$ follows the Wishart distribution (Wishart, 1928, Definition 3.4.1) with n degrees of freedom and scale matrix \mathbf{I}_{d_w} . Therefore, $(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w)^{-1} \sim \mathcal{W}^{-1}(\mathbf{I}_{d_w}, n)$ follows the inverse Wishart distribution (Mardia et al., 2024, §3.8), whose mean takes the form (Mardia et al., 2024, (3.8.3))

$$\mathbb{E}_{\tilde{\mathcal{S}}} \left[(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w)^\dagger \right] = \frac{1}{n - d_w - 1} \mathbf{I}_{d_w}.$$

Then, we have

$$\mathbb{E}_{\tilde{\mathcal{S}}} \left[\left(\mathbf{\Sigma}_w^{-1/2} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w \mathbf{\Sigma}_w^{-1/2} \right)^\dagger \right] = \frac{1}{n - d_w - 1} \mathbf{V}_w \mathbf{V}_w^\top.$$

Therefore, (8) implies

$$\begin{aligned} \text{Var}(f_{w2s}) &= \frac{\sigma^2}{N} \frac{1}{n - d_w - 1} \text{tr} \left(\mathbf{V}_w^\top \mathbb{E}_{\mathcal{S}_x} \left[\mathbf{\Sigma}_w^{-1/2} \mathbf{\Phi}_w^\top \mathbf{P}_s \mathbf{\Phi}_w \mathbf{\Sigma}_w^{-1/2} \right] \mathbf{V}_w \right) \\ &= \frac{\sigma^2}{N} \frac{1}{n - d_w - 1} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\mathbf{V}_w^\top \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{P}_s \mathbf{\Gamma}_w \mathbf{V}_w^\top \mathbf{V}_w \right] \right) \\ &= \frac{\sigma^2}{N} \frac{1}{n - d_w - 1} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\mathbf{\Gamma}_w^\top \mathbf{P}_s \mathbf{\Gamma}_w \right] \right). \end{aligned} \quad (10)$$

Recall that $\mathbf{P}_s = \mathbf{\Phi}_s \mathbf{\Phi}_s^\dagger$. Let $\mathbf{\Gamma}_s = \mathbf{\Gamma} \mathbf{V}_s \in \mathbb{R}^{N \times d_s}$, and we can write

$$\mathbf{P}_s = (\mathbf{\Phi}_s \mathbf{\Sigma}_s^{-1/2}) (\mathbf{\Phi}_s \mathbf{\Sigma}_s^{-1/2})^\dagger = (\mathbf{\Gamma}_s \mathbf{V}_s^\top) (\mathbf{\Gamma}_s \mathbf{V}_s^\top)^\dagger = \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger.$$

Therefore, with $\mathbf{\Gamma}_w = \mathbf{\Gamma} \mathbf{V}_w$ and $\mathbf{\Gamma}_s = \mathbf{\Gamma} \mathbf{V}_s$, we can decompose

$$\begin{aligned} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\mathbf{\Gamma}_w^\top \mathbf{P}_s \mathbf{\Gamma}_w \right] \right) &= \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{\Gamma}_w^\top \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger \mathbf{\Gamma}_w \right) \right] \\ &= \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{V}_w^\top \mathbf{V}_s \mathbf{V}_s^\top \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger \mathbf{\Gamma}_w \right) \right] \\ &\quad + \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger \mathbf{\Gamma}_w \right) \right]. \end{aligned}$$

For the first term, since $\mathbf{\Gamma}_w \mathbf{V}_w^\top \mathbf{V}_s = \mathbf{\Gamma} \mathbf{V}_w \mathbf{V}_w^\top \mathbf{V}_s$ and $\mathbf{\Gamma}_s = \mathbf{\Gamma} \mathbf{V}_s$, the range of $\mathbf{\Gamma}_w \mathbf{V}_w^\top \mathbf{V}_s$ is a subspace of that of $\mathbf{\Gamma}_s$ and therefore,

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{V}_w^\top \mathbf{V}_s \mathbf{V}_s^\top \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger \mathbf{\Gamma}_w \right) \right] &= \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{V}_s^\top \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger \mathbf{\Gamma}_w \mathbf{V}_w^\top \mathbf{V}_s \right) \right] \\ &= \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{V}_s^\top \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{\Gamma}_w \mathbf{V}_w^\top \mathbf{V}_s \right) \right] \\ &= \text{tr} \left(\mathbf{V}_s^\top \mathbf{V}_w \mathbb{E}_{\mathcal{S}_x} \left[\mathbf{\Gamma}_w^\top \mathbf{\Gamma}_w \right] \mathbf{V}_w^\top \mathbf{V}_s \right). \end{aligned}$$

Since $\mathbb{E}_{\mathcal{S}_x} \left[\mathbf{\Gamma}_w^\top \mathbf{\Gamma}_w \right] = N \mathbf{I}_{d_w}$, we have

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x} \left[\text{tr} \left(\mathbf{V}_w^\top \mathbf{V}_s \mathbf{V}_s^\top \mathbf{V}_w \mathbf{\Gamma}_w^\top \mathbf{\Gamma}_s \mathbf{\Gamma}_s^\dagger \mathbf{\Gamma}_w \right) \right] &= N \text{tr} \left(\mathbf{V}_s^\top \mathbf{V}_w \mathbf{V}_w^\top \mathbf{V}_s \right) \\ &= N \left\| \mathbf{V}_s^\top \mathbf{V}_w \right\|_F^2 \\ &= N d_{s \wedge w}. \end{aligned}$$

For the second term, we first observe that the row space of $\Gamma_w \mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top)$ is orthogonal to that of $\Gamma_s = \Gamma \mathbf{V}_s$, and therefore, $\Gamma_w \mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top)$ and Γ_s are independent, which implies

$$\mathbb{E}_{\mathcal{S}_x} [\text{tr} (\mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \Gamma_w^\top \Gamma_s \Gamma_s^\dagger \Gamma_w)] = \text{tr} (\mathbb{E} [\Gamma_w \mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \Gamma_w^\top] \mathbb{E} [\Gamma_s \Gamma_s^\dagger]).$$

Since Γ consists of independent isotropic rows, so do $\Gamma_s = \Gamma \mathbf{V}_s \in \mathbb{R}^{N \times d_s}$ and $\Gamma_w = \Gamma \mathbf{V}_w \in \mathbb{R}^{N \times d_w}$, which implies

$$\mathbb{E} [\Gamma_s \Gamma_s^\dagger] = \frac{d_s}{N} \mathbf{I}_N \quad \text{and} \quad \mathbb{E} [\Gamma_w^\top \Gamma_w] = N \mathbf{I}_{d_w}.$$

Then, we have

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x} [\text{tr} (\mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \Gamma_w^\top \Gamma_s \Gamma_s^\dagger \Gamma_w)] &= \text{tr} (\mathbb{E} [\Gamma_w \mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \Gamma_w^\top] \mathbb{E} [\Gamma_s \Gamma_s^\dagger]) \\ &= \frac{d_s}{N} \text{tr} (\mathbb{E} [\Gamma_w \mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \Gamma_w^\top]) \\ &= \frac{d_s}{N} \text{tr} (\mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w \mathbb{E} [\Gamma_w^\top \Gamma_w]) \\ &= \frac{d_s}{N} N \text{tr} (\mathbf{V}_w^\top (\mathbf{I}_d - \mathbf{V}_s \mathbf{V}_s^\top) \mathbf{V}_w) \\ &= d_s (d_w - d_{s \wedge w}). \end{aligned}$$

Combining the two terms, we have

$$\text{tr} (\mathbb{E}_{\mathcal{S}_x} [\Gamma_w^\top \mathbf{P}_s \Gamma_w]) = N d_{s \wedge w} + d_s (d_w - d_{s \wedge w}).$$

Then, by (10), the variance is exactly characterized by

$$\begin{aligned} \text{Var}(f_{w2s}) &= \frac{\sigma^2}{N} \frac{N d_{s \wedge w} + d_s (d_w - d_{s \wedge w})}{n - d_w - 1} \\ &= \frac{\sigma^2}{n - d_w - 1} \left(d_{s \wedge w} + \frac{d_s}{N} (d_w - d_{s \wedge w}) \right). \end{aligned}$$

- (b) **Sub-gaussian features:** Relaxing the Gaussian feature assumption, when $\tilde{\Gamma}_w$ consists of *i.i.d.* sub-gaussian random vectors that are zero-mean and isotropic (i.e. $\mathbb{E}[\tilde{\gamma}_i] = \mathbf{0}_{d_w}$ and $\mathbb{E}[\tilde{\gamma}_i \tilde{\gamma}_i^\top] = \mathbf{I}_{d_w}$), with $n \geq \Omega(d_w)$, Lemma B.3 implies that

$$\mathbb{E}_{\tilde{\mathcal{S}}} [(\tilde{\Gamma}_w^\top \tilde{\Gamma}_w)^\dagger] \preceq O\left(\frac{1}{n}\right) \mathbf{I}_{d_w},$$

and therefore,

$$\mathbb{E}_{\tilde{\mathcal{S}}} \left[\left(\Sigma_w^{-1/2} \tilde{\Phi}_w^\top \tilde{\Phi}_w \Sigma_w^{-1/2} \right)^\dagger \right] \preceq O\left(\frac{1}{n}\right) \mathbf{V}_w \mathbf{V}_w^\top.$$

Then, via an analogous argument as (10), (8) implies that

$$\text{Var}(f_{w2s}) \leq \frac{\sigma^2}{N} O\left(\frac{1}{n}\right) \text{tr} (\mathbb{E}_{\mathcal{S}_x} [\Gamma_w^\top \mathbf{P}_s \Gamma_w]). \quad (11)$$

We observe that in the analysis of the Gaussian feature case, the characterization

$$\text{tr} (\mathbb{E}_{\mathcal{S}_x} [\Gamma_w^\top \mathbf{P}_s \Gamma_w]) = (N - d_s) d_{s \wedge w} + d_s d_w$$

does not involve the Gaussianity of Γ and therefore holds for general subgaussian features. This leads to an upper bound on the variance:

$$\begin{aligned} \text{Var}(f_{w2s}) &\leq \frac{\sigma^2}{N} O\left(\frac{1}{n}\right) (N d_{s \wedge w} + d_s (d_w - d_{s \wedge w})) \\ &\lesssim \frac{\sigma^2}{n} \left(d_{s \wedge w} + \frac{d_s}{N} (d_w - d_{s \wedge w}) \right). \end{aligned}$$

□

Lemma B.3 (Adapting (Vershynin, 2010) Theorem 5.39). *Let $\tilde{\Gamma}_w = [\tilde{\gamma}_1, \dots, \tilde{\gamma}_n]^\top$ be an $n \times d_w$ matrix whose rows $\tilde{\gamma}_1, \dots, \tilde{\gamma}_n$ consist of i.i.d. sub-gaussian random vectors that are zero-mean and isotropic (i.e. $\mathbb{E}[\tilde{\gamma}_i] = \mathbf{0}_{d_w}$ and $\mathbb{E}[\tilde{\gamma}_i \tilde{\gamma}_i^\top] = \mathbf{I}_{d_w}$). When $n \geq \Omega(d_w)$, we have*

$$\mathbb{E} \left[\left\| \left(\tilde{\Gamma}_w^\top \tilde{\Gamma}_w \right)^\dagger \right\|_2 \right] \leq O \left(\frac{1}{n} \right),$$

where $\Omega(\cdot)$ and $O(\cdot)$ suppresses constants that depend only on the sub-gaussian norm $\|\tilde{\gamma}_i\|_{\psi_2} = \sup_{\mathbf{v} \in \mathbb{S}^{d_w-1}} \sup_{p \geq 1} (\mathbb{E}[|\tilde{\gamma}_i^\top \mathbf{v}|^p])^{1/p} / \sqrt{p}$, independent of n, d_w .

Proof of Lemma B.3. Let $\sigma_{\min}(\tilde{\Gamma}_w^\top \tilde{\Gamma}_w)$ be the smallest singular value of $\tilde{\Gamma}_w^\top \tilde{\Gamma}_w$. Leveraging (Vershynin, 2010) Theorem 5.39, we notice that for $n \geq \Omega(d_w)$, there exist constants $c_1, c_2 > 0$ that depend only on the sub-gaussian norm $\|\tilde{\gamma}_i\|_{\psi_2}$ such that

$$\Pr \left[\sigma_{\min}(\tilde{\Gamma}_w^\top \tilde{\Gamma}_w) < \left(\sqrt{n} - c_1 \sqrt{d_w} - t \right)^2 \right] \leq \exp(-c_2 t^2).$$

Therefore, we have

$$\Pr \left[\frac{1}{\sigma_{\min}(\tilde{\Gamma}_w^\top \tilde{\Gamma}_w)} > t \right] \leq \exp \left(-c_2 \left(\sqrt{n} - c_1 \sqrt{d_w} - \sqrt{\frac{1}{t}} \right)^2 \right).$$

Notice that for any non-negative random variable Z with a cumulative density function $F_Z(z)$,

$$\begin{aligned} \mathbb{E}[Z] &= \int_0^\infty z dF_Z(z) = - \int_0^\infty z d(1 - F_Z(z)) \\ &= [z(1 - F_Z(z))]_0^\infty + \int_0^\infty (1 - F_Z(z)) dz \\ &= \int_0^\infty \Pr[Z > z] dz. \end{aligned}$$

Therefore, we have

$$\mathbb{E} \left[\frac{1}{\sigma_{\min}(\tilde{\Gamma}_w^\top \tilde{\Gamma}_w)} \right] \leq \int_0^\infty \exp \left(-c_2 \left(\sqrt{n} - c_1 \sqrt{d_w} - \sqrt{\frac{1}{t}} \right)^2 \right) dt.$$

Let $t_0 = 1 / (\sqrt{n} - c_1 \sqrt{d_w})^2$ such that $\sqrt{n} - c_1 \sqrt{d_w} - \sqrt{\frac{1}{t}} = 0$ and

$$\int_0^{t_0} \exp \left(-c_2 \left(\sqrt{n} - c_1 \sqrt{d_w} - \sqrt{\frac{1}{t}} \right)^2 \right) dt \leq t_0$$

Then, we have

$$\begin{aligned}
 \mathbb{E} \left[\frac{1}{\sigma_{\min}(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w)} \right] &\leq \int_0^\infty \exp \left(-c_2 \left(\sqrt{n} - c_1 \sqrt{d_w} - \sqrt{\frac{1}{t}} \right)^2 \right) dt \\
 &\leq t_0 + \int_{t_0}^\infty \exp \left(-c_2 \left(\sqrt{n} - c_1 \sqrt{d_w} - \sqrt{\frac{1}{t}} \right)^2 \right) dt \\
 &= t_0 + 2 \int_0^{\sqrt{n} - c_1 \sqrt{d_w}} \exp(-c_2 u^2) \left(\sqrt{n} - c_1 \sqrt{d_w} - u \right)^{-3} du \\
 &= t_0 + \frac{2}{(\sqrt{n} - c_1 \sqrt{d_w})^2} \int_0^1 \exp \left(-c_2 \left(\sqrt{n} - c_1 \sqrt{d_w} \right)^2 u^2 \right) (1-u)^{-3} du \\
 &= \frac{1}{(\sqrt{n} - c_1 \sqrt{d_w})^2} + \frac{2}{(\sqrt{n} - c_1 \sqrt{d_w})^2} \left(\int_0^1 \exp(-\Omega(u^2)) (1-u)^{-3} du \right) \\
 &= O \left(\frac{1}{(\sqrt{n} - c_1 \sqrt{d_w})^2} \right).
 \end{aligned}$$

When $n \geq \Omega(d_w)$, we have $\sqrt{n} - c_1 \sqrt{d_w} \geq \Omega(\sqrt{n})$, and therefore ,

$$\mathbb{E} \left[\left\| \left(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w \right)^\dagger \right\|_2 \right] \leq \mathbb{E} \left[\frac{1}{\sigma_{\min}(\tilde{\mathbf{\Gamma}}_w^\top \tilde{\mathbf{\Gamma}}_w)} \right] \leq O \left(\frac{1}{n} \right).$$

□

B.2. Proof of Proposition 3.3 and Corollary 3.4

Proof of Proposition 3.3 and Corollary 3.4. The excess risk of the finetuned weak teacher $f_w(\mathbf{x}) = \phi_w(\mathbf{x})^\top \boldsymbol{\theta}_w$ can be expressed as

$$\begin{aligned}
 \mathbf{ER}(f_w) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\mathbb{E}_{f_w} [(f_w(\mathbf{x}) - f_*(\mathbf{x}))^2]] \\
 &= \mathbb{E}_{\tilde{\mathcal{S}}} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Phi}}_w \boldsymbol{\theta}_w - \tilde{\mathbf{f}}_* \right\|_2^2 \right],
 \end{aligned}$$

where $\tilde{\mathbf{f}}_* = [\mathbf{f}_*(\tilde{\mathbf{x}}_1), \dots, \mathbf{f}_*(\tilde{\mathbf{x}}_n)]^\top \in \mathbb{R}^n$; and we recall that $\tilde{\mathbf{\Phi}}_w = [\phi_w(\tilde{\mathbf{x}}_1), \dots, \phi_w(\tilde{\mathbf{x}}_n)]^\top$. Notice that the randomness of $\boldsymbol{\theta}_w$ comes from the SFT samples $\tilde{\mathcal{S}} \sim \mathcal{D}(f_*)^n$.

Observe that the solution of (2) as $\alpha_w \rightarrow 0$ is given by $\boldsymbol{\theta}_w = \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{y}}$, where $\tilde{\mathbf{y}} = \tilde{\mathbf{f}}_* + \tilde{\mathbf{z}}$ is the noisy label vector with $\tilde{\mathbf{z}} \sim \mathcal{N}(\mathbf{0}_n, \sigma^2 \mathbf{I}_n)$. Therefore, with the randomness over $\tilde{\mathcal{S}} \sim \mathcal{D}(f_*)^n$, we have

$$\begin{aligned}
 \mathbf{ER}(f_w) &= \mathbb{E} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Phi}}_w \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{y}} - \tilde{\mathbf{f}}_* \right\|_2^2 \right] \\
 &= \mathbb{E} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Phi}}_w \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{z}} + \left(\tilde{\mathbf{\Phi}}_w \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{f}}_* - \tilde{\mathbf{f}}_* \right) \right\|_2^2 \right] \\
 &= \underbrace{\mathbb{E} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Phi}}_w \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{z}} \right\|_2^2 \right]}_{\mathbf{Var}(f_w)} + \underbrace{\mathbb{E} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Phi}}_w \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{f}}_* - \tilde{\mathbf{f}}_* \right\|_2^2 \right]}_{\mathbf{Bias}(f_w)}.
 \end{aligned}$$

For bias, by the definition of finetuning capacity (see Definition 2.2), we have

$$\mathbf{Bias}(f_w) = \frac{1}{n} \mathbb{E} \left[\left\| \tilde{\mathbf{\Phi}}_w \tilde{\mathbf{\Phi}}_w^\dagger \tilde{\mathbf{f}}_* - \tilde{\mathbf{f}}_* \right\|_2^2 \right] = \frac{\rho_w(n)}{n}.$$

We observe that $\text{Bias}(f_w) \leq \rho_w$ by Lemma B.1. Notice that Lemma B.1 also implies $\text{Bias}(f_s) = \rho_s(n)/n \leq \rho_s$.

For variance, we observe that

$$\begin{aligned} \text{Var}(f_w) &= \frac{1}{n} \mathbb{E} \left[\left\| \tilde{\Phi}_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{z}} \right\|_2^2 \right] \\ &= \frac{1}{n} \mathbb{E} \left[\text{tr} \left(\tilde{\Phi}_w \tilde{\Phi}_w^\dagger \tilde{\mathbf{z}} \tilde{\mathbf{z}}^\top \right) \right] \\ &= \frac{\sigma^2}{n} \mathbb{E} \left[\text{tr} \left(\tilde{\Phi}_w \tilde{\Phi}_w^\dagger \right) \right]. \end{aligned}$$

By Assumption 2.6, since $\text{rank}(\tilde{\Phi}_w) = d_w$ almost surely, we have

$$\text{Var}(f_w) = \frac{\sigma^2}{n} \mathbb{E} \left[\text{tr} \left(\tilde{\Phi}_w \tilde{\Phi}_w^\dagger \right) \right] = \frac{\sigma^2 d_w}{n}.$$

□

B.3. Proof of Proposition 3.5

Proof of Proposition 3.5. Noticing that with $\text{rank}(\tilde{\Phi}_w) = d_w$ and $\text{rank}(\tilde{\Phi}_s) = \text{rank}(\Phi_s) = d_s$ almost surely, the excess risks of f_w, f_s, f_c are characterized exactly in Proposition 3.3 and Corollary 3.4, and $\text{ER}(f_{w2s})$ is upper bounded by Theorem 3.1. Therefore, by directly plugging in the excess risks to the definitions of PGR and OPR, we have

$$\begin{aligned} \text{PGR} &= \frac{\text{ER}(f_w) - \text{ER}(f_{w2s})}{\text{ER}(f_w) - \text{ER}(f_c)} \\ &\geq \left(\sigma^2 \frac{d_w}{n} + \frac{\rho_w(n)}{n} - \frac{\sigma^2}{n - d_w - 1} \left(d_{s \wedge w} + \frac{d_s}{N} (d_w - d_{s \wedge w}) \right) - \left(\frac{\rho_w(n)}{n} + \frac{\rho_s(N)}{N} \right) \right) \\ &\quad \left(\sigma^2 \frac{d_w}{n} + \frac{\rho_w(n)}{n} - \sigma^2 \frac{d_s}{N + n} - \frac{\rho_s(N + n)}{N + n} \right)^{-1} \\ &\geq \left(\sigma^2 \frac{d_w}{n} - \sigma^2 \frac{d_{s \wedge w} + (d_w - d_{s \wedge w}) d_s / N}{n - d_w - 1} - \frac{\rho_s(N)}{N} \right) / \left(\sigma^2 \frac{d_w}{n} + \frac{\rho_w(n)}{n} \right), \\ &\geq \left(\sigma^2 \frac{d_w}{n} - \sigma^2 \frac{d_{s \wedge w} + (d_w - d_{s \wedge w}) d_s / N}{n - d_w - 1} - \rho_s \right) / \left(\sigma^2 \frac{d_w}{n} + \rho_w \right), \end{aligned} \tag{12}$$

and

$$\begin{aligned} \text{OPR} &= \frac{\text{ER}(f_s)}{\text{ER}(f_{w2s})} \\ &\geq \left(\sigma^2 \frac{d_s}{n} + \frac{\rho_s(n)}{n} \right) / \left(\sigma^2 \frac{d_{s \wedge w} + (d_w - d_{s \wedge w}) d_s / N}{n - d_w - 1} + \left(\frac{\rho_w(n)}{n} + \frac{\rho_s(N)}{N} \right) \right) \\ &\geq \sigma^2 \frac{d_s}{n} / \left(\sigma^2 \frac{d_{s \wedge w} + (d_w - d_{s \wedge w}) d_s / N}{n - d_w - 1} + \rho_w + \rho_s \right). \end{aligned} \tag{13}$$

When taking $n = d_w + q + 1$ for some small constant $q \in \mathbb{N}$, we observe that

$$\begin{aligned} \text{PGR} &\geq \left(\sigma^2 \frac{d_w}{n} - \sigma^2 \frac{d_{s \wedge w} + (d_w - d_{s \wedge w}) d_s / N}{n - d_w - 1} - \rho_s \right) / \left(\sigma^2 \frac{d_w}{n} + \rho_w \right) \\ &\geq \left(\frac{d_w}{d_w + q + 1} - \frac{d_{s \wedge w}}{q} - \frac{d_s}{N} \frac{d_w - d_{s \wedge w}}{q} - \frac{\rho_s}{\sigma^2} \right) / \left(\frac{d_w}{d_w + q + 1} + \frac{\rho_w}{\sigma^2} \right) \\ &\geq \left(\frac{d_w}{d_w + q + 1} - \frac{d_{s \wedge w}}{q} - \frac{d_s}{N} \frac{d_w - d_{s \wedge w}}{q} - \frac{\rho_s}{\sigma^2} - \frac{\rho_w}{\sigma^2} \right) / \left(\frac{d_w}{d_w + q + 1} + \frac{\rho_w}{\sigma^2} - \frac{\rho_w}{\sigma^2} \right) \\ &= 1 - \frac{n}{d_w} \left(\frac{d_{s \wedge w}}{q} + \frac{d_s}{N} \frac{d_w - d_{s \wedge w}}{q} + \frac{\rho_w + \rho_s}{\sigma^2} \right) \\ &= 1 - \frac{n}{q} \frac{d_{s \wedge w} + (d_w - d_{s \wedge w}) d_s / N}{d_w} - \frac{n}{d_w} \frac{\rho_w + \rho_s}{\sigma^2}, \end{aligned}$$

and

$$\begin{aligned}
 \text{OPR} &\geq \sigma^2 \frac{d_s}{n} / \left(\sigma^2 \frac{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}{n - d_w - 1} + \rho_w + \rho_s \right) \\
 &= \frac{d_s}{n} / \left(\frac{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}{q} + \frac{\rho_w + \rho_s}{\sigma^2} \right) \\
 &= \left(\frac{n}{q} \frac{d_{s \wedge w} + (d_w - d_{s \wedge w})d_s/N}{d_s} + \frac{n}{d_s} \frac{\rho_w + \rho_s}{\sigma^2} \right)^{-1}.
 \end{aligned}$$

□

B.4. Proof of Corollary 3.6

Proof of Corollary 3.6. Recall the notations introduced for conciseness:

$$d_{w2s}(N) = d_{s \wedge w} + (d_w - d_{s \wedge w}) \frac{d_s}{N}, \quad \varrho = \frac{\rho_w + \rho_s}{\sigma^2}.$$

Then, the lower bounds for **PGR** and **OPR** in Proposition 3.5 can be expressed in terms of $d_{w2s}(N)$ and ϱ as

$$\text{PGR} \geq 1 - \frac{d_{w2s}(N)}{d_w} - \frac{d_w + 1}{d_w} \varrho - \frac{d_w + 1}{d_w} \frac{d_{w2s}(N)}{q} - q \frac{\varrho}{d_w},$$

and

$$\text{OPR} \geq \left(\frac{d_{w2s}(N)}{d_s} + \frac{d_w + 1}{d_s} \varrho + \frac{d_{w2s}(N)}{d_s} \frac{d_w + 1}{q} + q \frac{\varrho}{d_s} \right)^{-1}.$$

Both lower bounds are maximized when the last two terms in the expressions that involve q are minimized, which is achieved when $q = \sqrt{(d_w + 1) d_{w2s}(N) / \varrho}$. Substituting the optimal q back into the expressions yields the best lower bounds for **PGR** and **OPR**:

$$\begin{aligned}
 \text{PGR} &\geq 1 - \frac{d_{w2s}(N)}{d_w} - \varrho \frac{d_w + 1}{d_w} - 2 \sqrt{\varrho \frac{d_w + 1}{d_w} \frac{d_{w2s}(N)}{d_w}} \\
 &= 1 - \left(\sqrt{\frac{d_{w2s}(N)}{d_w}} + \sqrt{\varrho \frac{d_w + 1}{d_w}} \right)^2,
 \end{aligned}$$

and

$$\begin{aligned}
 \text{OPR} &\geq \left(\frac{d_{w2s}(N)}{d_s} + \varrho \frac{d_w + 1}{d_s} + 2 \sqrt{\varrho \frac{d_w + 1}{d_s} \frac{d_{w2s}(N)}{d_s}} \right)^{-1} \\
 &= \left(\sqrt{\frac{d_{w2s}(N)}{d_s}} + \sqrt{\varrho \frac{d_w + 1}{d_s}} \right)^{-2}.
 \end{aligned}$$

□

C. Ridge regression analysis

In this section, we investigate the more realistic scenario where the weak and strong feature covariances are not exactly low-rank but admit a small number of dominating eigenvalues.

Concretely, we consider the same data distribution $(\mathbf{x}, y) \sim \mathcal{D}(f_*)$ with $y = f_*(\mathbf{x}) + z$ for some independent Gaussian label noise $z \sim \mathcal{N}(0, \sigma^2)$ and an unknown ground truth function $f_* : \mathcal{X} \rightarrow \mathbb{R}$ as in Section 2. Under the same sub-gaussian feature assumption as in Assumption 2.1, we adapt Definitions 2.3 and 2.4 to the ridge regression setting as follows.

Assumption C.1 (Data distribution). Let $\phi_s : \mathcal{X} \rightarrow \mathbb{R}^d$ and $\phi_w : \mathcal{X} \rightarrow \mathbb{R}^d$ be the strong and weak pretrained models that take $\mathbf{x} \sim \mathcal{D}$ and output pretrained features $\phi_s(\mathbf{x}), \phi_w(\mathbf{x}) \in \mathbb{R}^d$, respectively.

- (i) **Ground truth:** Assume f_* can be expressed as a linear function over an unknown ground truth feature $\phi_* : \mathcal{X} \rightarrow \mathbb{R}^d$ such that $f_*(\cdot) = \phi_*(\cdot)^\top \boldsymbol{\theta}_*$ for some fixed $\boldsymbol{\theta}_* \in \mathbb{R}^d$.
- (ii) **Sub-gaussian features** (Assumption 2.1): Let $\phi_w(\mathbf{x}), \phi_s(\mathbf{x}), \phi_*(\mathbf{x})$ be zero-mean sub-gaussian random vectors with $\mathbb{E}[\phi_w(\mathbf{x})] = \mathbb{E}[\phi_s(\mathbf{x})] = \mathbb{E}[\phi_*(\mathbf{x})] = \mathbf{0}_d$, and

$$\mathbb{E}[\phi_w(\mathbf{x})\phi_w(\mathbf{x})^\top] = \boldsymbol{\Sigma}_w, \quad \mathbb{E}[\phi_s(\mathbf{x})\phi_s(\mathbf{x})^\top] = \boldsymbol{\Sigma}_s, \quad \mathbb{E}[\phi_*(\mathbf{x})\phi_*(\mathbf{x})^\top] = \boldsymbol{\Sigma}_*.$$

For conciseness, we assume without loss of generality that these features are roughly normalized, *i.e.*, $\|\boldsymbol{\Sigma}_w\|_2 \asymp 1$, $\|\boldsymbol{\Sigma}_s\|_2 \asymp 1$, and $\|\boldsymbol{\Sigma}_*\|_2 \asymp 1$.

- (iii) **Low intrinsic dimension:** Let $\boldsymbol{\Sigma}_s$ and $\boldsymbol{\Sigma}_w$ both be **positive-definite** with spectral decompositions $\boldsymbol{\Sigma}_s = \mathbf{V}_s \boldsymbol{\Lambda}_s \mathbf{V}_s^\top$ and $\boldsymbol{\Sigma}_w = \mathbf{V}_w \boldsymbol{\Lambda}_w \mathbf{V}_w^\top$, where $\boldsymbol{\Lambda}_s, \boldsymbol{\Lambda}_w \in \mathbb{R}^{d \times d}$ are diagonal matrices with positive eigenvalues in decreasing order; while $\mathbf{V}_s \in \mathbb{R}^{d \times d}$ and $\mathbf{V}_w \in \mathbb{R}^{d \times d}$ are orthogonal matrices consisting of the corresponding orthonormal eigenvectors. The low intrinsic dimension of FT implies that $\boldsymbol{\Lambda}_s = \text{diag}(\lambda_1^s, \dots, \lambda_d^s)$ and $\boldsymbol{\Lambda}_w = \text{diag}(\lambda_1^w, \dots, \lambda_d^w)$ consist of a few dominating eigenvalues, while the rest of the eigenvalues are negligible, *i.e.*, there exist $d_s, d_w \ll d$ such that $\sum_{i>d_s} \lambda_i^s \ll \text{tr}(\boldsymbol{\Sigma}_s)$ and $\sum_{i>d_w} \lambda_i^w \ll \text{tr}(\boldsymbol{\Sigma}_w)$. Here,

$$\text{tr}(\boldsymbol{\Sigma}_s) \lesssim d_s \quad \text{and} \quad \text{tr}(\boldsymbol{\Sigma}_w) \lesssim d_w$$

effectively measure the intrinsic dimensions of ϕ_s and ϕ_w .

Remark C.2 (Weak-strong similarity). In place of correlation dimension (Definition 2.4) in the ridgeless setting, for the ridge regression analysis, we measure the similarity between the weak and strong models directly through $\text{tr}(\boldsymbol{\Sigma}_s \boldsymbol{\Sigma}_w)$. Notice that

$$\text{tr}(\boldsymbol{\Sigma}_s \boldsymbol{\Sigma}_w) \leq \min \{ \text{tr}(\boldsymbol{\Sigma}_s) \|\boldsymbol{\Sigma}_w\|_2, \text{tr}(\boldsymbol{\Sigma}_w) \|\boldsymbol{\Sigma}_s\|_2 \} \lesssim \min \{ \text{tr}(\boldsymbol{\Sigma}_s), \text{tr}(\boldsymbol{\Sigma}_w) \}.$$

In particular, when $\boldsymbol{\Sigma}_s$ and $\boldsymbol{\Sigma}_w$ admit low intrinsic dimensions, $\text{tr}(\boldsymbol{\Sigma}_s \boldsymbol{\Sigma}_w)$ can be much smaller than $\min \{ \text{tr}(\boldsymbol{\Sigma}_s), \text{tr}(\boldsymbol{\Sigma}_w) \}$ if their eigenvectors corresponding to the dominating eigenvalues are almost orthogonal.

Remark C.3 (FT approximation errors). It is worth noting that under the ground truth and positive-definite covariance assumptions in Assumption C.1(i, iii), the FT approximation errors in Definition 2.2 satisfy

$$\begin{aligned} \rho_s &= \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_s(\mathbf{x})^\top \boldsymbol{\theta} - f_*(\mathbf{x}))^2] = 0 \quad (\text{when } \boldsymbol{\theta} = \boldsymbol{\Sigma}_s^{-1} \boldsymbol{\Sigma}_* \boldsymbol{\theta}_*), \\ \rho_w &= \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(\phi_w(\mathbf{x})^\top \boldsymbol{\theta} - f_*(\mathbf{x}))^2] = 0 \quad (\text{when } \boldsymbol{\theta} = \boldsymbol{\Sigma}_w^{-1} \boldsymbol{\Sigma}_* \boldsymbol{\theta}_*). \end{aligned} \tag{14}$$

In place of Definition 2.2, with positive-definite covariances in Assumption C.1, we measure the alignment between the ground truth feature ϕ_* and the weak/strong feature ϕ_w, ϕ_s through

$$\varrho_s = \|\boldsymbol{\Sigma}_s^{-1/2} \boldsymbol{\Sigma}_*^{1/2} \boldsymbol{\theta}_*\|_2^2, \quad \varrho_w = \|\boldsymbol{\Sigma}_w^{-1/2} \boldsymbol{\Sigma}_*^{1/2} \boldsymbol{\theta}_*\|_2^2.$$

Intuitively, for $\boldsymbol{\Sigma}_s$ and $\boldsymbol{\Sigma}_w$ with a few dominating eigenvalues (Assumption C.1(iii)), ϱ_s and ϱ_w are small if the eigensubspace associated with non-negligible eigenvalues of $\boldsymbol{\Sigma}_*$ is fully covered by the eigensubspaces associated with the dominating eigenvalues of $\boldsymbol{\Sigma}_s$ and $\boldsymbol{\Sigma}_w$, respectively.

The W2S FT under ridge regression consists of two steps.

- (a) First, the weak teacher $f_w(\mathbf{x}) = \phi_w(\mathbf{x})^\top \boldsymbol{\theta}_w$ is supervisedly finetuned over $\tilde{\mathcal{S}}$:

$$\boldsymbol{\theta}_w = \underset{\boldsymbol{\theta} \in \mathbb{R}^d}{\text{argmin}} \frac{1}{n} \|\tilde{\boldsymbol{\Phi}}_w \boldsymbol{\theta} - \tilde{\mathbf{y}}\|_2^2 + \alpha_w \|\boldsymbol{\theta}\|_2^2, \quad \alpha_w > 0. \tag{15}$$

- (b) Then, the W2S model $f_{w2s}(\mathbf{x}) = \phi_s(\mathbf{x})^\top \boldsymbol{\theta}_{w2s}$ is finetuned over the strong feature ϕ_s through the unlabeled samples \mathcal{S}_x and their pseudo-labels generated by the weak teacher model:

$$\boldsymbol{\theta}_{w2s} = \underset{\boldsymbol{\theta} \in \mathbb{R}^d}{\text{argmin}} \frac{1}{N} \|\boldsymbol{\Phi}_s \boldsymbol{\theta} - \boldsymbol{\Phi}_w \boldsymbol{\theta}_w\|_2^2 + \alpha_{w2s} \|\boldsymbol{\theta}\|_2^2, \quad \alpha_{w2s} > 0. \tag{16}$$

Theorem C.4 (W2S under ridge regression). *Let $\varrho_w = \left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2$ and $\varrho_s = \left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2$. Under Assumption C.1, the generalization error of W2S FT via ridge regression with fixed $\alpha_w, \alpha_{w2s} > 0$, $\mathbf{ER}(f_{w2s}) = \mathbf{Var}(f_{w2s}) + \mathbf{Bias}(f_{w2s})$, is upper bounded by*

$$\mathbf{Var}(f_{w2s}) \leq \frac{\sigma^2 \operatorname{tr}(\Sigma_s \Sigma_w)}{4(\alpha_w n)(\alpha_{w2s} N)}, \quad \mathbf{Bias}(f_{w2s}) \leq \alpha_w \varrho_w + \alpha_{w2s} \varrho_s.$$

In particular, when taking

$$\alpha_w = \left(\frac{\sigma^2 \operatorname{tr}(\Sigma_s \Sigma_w)}{4nN} \frac{\varrho_s}{\varrho_w^2} \right)^{1/3}, \quad \alpha_{w2s} = \left(\frac{\sigma^2 \operatorname{tr}(\Sigma_s \Sigma_w)}{4nN} \frac{\varrho_w}{\varrho_s^2} \right)^{1/3},$$

the excess risk of W2S FT is upper bounded by

$$\mathbf{ER}(f_{w2s}) \leq 3 \left(\frac{\sigma^2 \operatorname{tr}(\Sigma_s \Sigma_w)}{4nN} \varrho_s \varrho_w \right)^{1/3}.$$

Theorem C.4 conveys a similar high-level intuition as in Theorem 3.1 regarding the effect of the weak-strong similarity on the generalization error of W2S FT. In particular, the larger discrepancy between ϕ_s and ϕ_w (corresponding to the smaller $\operatorname{tr}(\Sigma_s \Sigma_w)$) leads to lower variance and therefore better W2S generalization.

Meanwhile, a key difference in W2S between the ridge and ridgeless settings (Theorem C.4 versus Theorem 3.1) is that the FT approximation errors in Theorem C.4, reflected by $\varrho_s = \left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2$ and $\varrho_w = \left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2$, can be compensated by larger sample sizes n, N and directly affect the sample complexity:

$$nN \asymp \sigma^2 \operatorname{tr}(\Sigma_s \Sigma_w) \varrho_s \varrho_w.$$

Such difference is a result of optimizing the regularization hyperparameters α_w, α_{w2s} in ridge regression that control the variance-bias tradeoff.

Proof of Theorem C.4. We first formalize some useful facts on the features and labels as in (9). In particular, the sub-gaussian assumption in Assumption C.1(ii) implies that for each $\mathbf{x} \sim \mathcal{D}$, the corresponding strong/weak feature $\phi_s(\mathbf{x}), \phi_w(\mathbf{x}) \in \mathbb{R}^d$, and the ground truth $f_*(\mathbf{x}) \in \mathbb{R}$ are simultaneously characterized by an independent subgaussian random vector $\gamma \in \mathbb{R}^d$ with $\mathbb{E}[\gamma] = \mathbf{0}_d$ and $\mathbb{E}[\gamma \gamma^\top] = \mathbf{I}_d$, i.e.,

$$\phi_s(\mathbf{x}) = \Sigma_s^{1/2} \gamma, \quad \phi_w(\mathbf{x}) = \Sigma_w^{1/2} \gamma, \quad f_*(\mathbf{x}) = \phi_*(\mathbf{x})^\top \theta_* = \gamma^\top \Sigma_*^{1/2} \theta_*.$$

Then, for \mathcal{S} and $\tilde{\mathcal{S}}$, there exist independent random matrices $\Gamma = [\gamma_1, \dots, \gamma_N]^\top \in \mathbb{R}^{N \times d}$ and $\tilde{\Gamma} = [\tilde{\gamma}_1, \dots, \tilde{\gamma}_n]^\top \in \mathbb{R}^{n \times d}$ consisting of i.i.d. zero-mean isotropic rows such that

$$\begin{aligned} \Phi_s &= \Gamma \Sigma_s^{1/2} = \Gamma_s \Lambda_s^{1/2} \mathbf{V}_s^\top, \\ \Phi_w &= \Gamma \Sigma_w^{1/2} = \Gamma_w \Lambda_w^{1/2} \mathbf{V}_w^\top, \\ \mathbf{y} &= \mathbf{f}_* + \mathbf{z}, \quad \mathbf{f}_* = \Gamma \Sigma_*^{1/2} \theta_*, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}_N, \sigma^2 \mathbf{I}_N), \\ \tilde{\Phi}_w &= \tilde{\Gamma} \Sigma_w^{1/2} = \tilde{\Gamma}_w \Lambda_w^{1/2} \mathbf{V}_w^\top, \\ \tilde{\mathbf{y}} &= \tilde{\mathbf{f}}_* + \tilde{\mathbf{z}}, \quad \tilde{\mathbf{f}}_* = \tilde{\Gamma} \Sigma_*^{1/2} \theta_*, \quad \tilde{\mathbf{z}} \sim \mathcal{N}(\mathbf{0}_n, \sigma^2 \mathbf{I}_n), \end{aligned} \tag{17}$$

where $\Gamma_s = \Gamma \mathbf{V}_s$, $\Gamma_w = \Gamma \mathbf{V}_w$, and $\tilde{\Gamma}_w = \tilde{\Gamma} \mathbf{V}_w$.

Variance-bias decomposition. Recall that the excess risk of W2S generalization $\mathbf{ER}(f_{w2s})$ can be decomposed into the variance and bias terms:

$$\begin{aligned} \mathbf{Var}(f_{w2s}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[(f_{w2s}(\mathbf{x}) - \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} [f_{w2s}(\mathbf{x})])^2 \right] \right], \\ \mathbf{Bias}(f) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[(\mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} [f_{w2s}(\mathbf{x})] - f_*(\mathbf{x}))^2 \right]. \end{aligned}$$

With $\alpha_w > 0$, (15) yields a weak teacher model $f_w(\mathbf{x}) = \phi_w(\mathbf{x})^\top \boldsymbol{\theta}_w$ with

$$\boldsymbol{\theta}_w = \left(\tilde{\boldsymbol{\Phi}}_w^\top \tilde{\boldsymbol{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\boldsymbol{\Phi}}_w^\top (\tilde{\mathbf{f}}_8 + \tilde{\mathbf{z}}).$$

Then, the W2S model $f_{w2s}(\mathbf{x}) = \phi_s(\mathbf{x})^\top \boldsymbol{\theta}_{w2s}$ is given by (16) with $\alpha_{w2s} > 0$:

$$\begin{aligned} \boldsymbol{\theta}_{w2s} &= (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_w \boldsymbol{\theta}_w \\ &= (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_w \left(\tilde{\boldsymbol{\Phi}}_w^\top \tilde{\boldsymbol{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\boldsymbol{\Phi}}_w^\top (\tilde{\mathbf{f}}_* + \tilde{\mathbf{z}}), \end{aligned}$$

which implies

$$\mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}}[\boldsymbol{\theta}_{w2s}] = (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_w \left(\tilde{\boldsymbol{\Phi}}_w^\top \tilde{\boldsymbol{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\boldsymbol{\Phi}}_w^\top \tilde{\mathbf{f}}_*.$$

Then, we can concretize the variance and bias terms as:

$$\begin{aligned} \text{Var}(f_{w2s}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[(f_{w2s}(\mathbf{x}) - \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}}[f_{w2s}(\mathbf{x})])^2 \right] \right] \\ &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \boldsymbol{\Sigma}_s^{1/2} (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_w \left(\tilde{\boldsymbol{\Phi}}_w^\top \tilde{\boldsymbol{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\boldsymbol{\Phi}}_w^\top \tilde{\mathbf{z}} \right\|_2^2 \right], \end{aligned} \quad (18)$$

and

$$\begin{aligned} \text{Bias}(f_{w2s}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[(\mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}}[f_{w2s}(\mathbf{x})] - f_*(\mathbf{x}))^2 \right] \\ &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \boldsymbol{\Phi}_s (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_w \left(\tilde{\boldsymbol{\Phi}}_w^\top \tilde{\boldsymbol{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\boldsymbol{\Phi}}_w^\top \tilde{\mathbf{f}}_* - \mathbf{f}_* \right\|_2^2 \right]. \end{aligned} \quad (19)$$

Now, we are ready to upper-bound the variance and bias terms separately.

Variance. Denote $\boldsymbol{\zeta} = \boldsymbol{\Lambda}_w^{1/2} \mathbf{V}_w^\top \left(\tilde{\boldsymbol{\Phi}}_w^\top \tilde{\boldsymbol{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\boldsymbol{\Phi}}_w^\top \tilde{\mathbf{z}} \in \mathbb{R}^d$, whose randomness comes from $\tilde{\mathcal{S}}$ only, independent of \mathcal{S}_x . Then, the variance term (18) can be expressed as

$$\begin{aligned} \text{Var}(f_{w2s}) &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \boldsymbol{\Sigma}_s^{1/2} (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_w \boldsymbol{\zeta} \right\|_2^2 \right] \\ &= \text{tr} \left(\mathbb{E}_{\mathcal{S}_s} \left(\boldsymbol{\Gamma}_w^\top \boldsymbol{\Gamma}_s (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Sigma}_s (\boldsymbol{\Phi}_s^\top \boldsymbol{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d)^{-1} \boldsymbol{\Phi}_s^\top \boldsymbol{\Gamma}_w \right) \mathbb{E}_{\tilde{\mathcal{S}}} [\boldsymbol{\zeta} \boldsymbol{\zeta}^\top] \right) \\ &= \text{tr} \left(\mathbb{E}_{\mathcal{S}_s} \left(\boldsymbol{\Gamma}_w^\top \boldsymbol{\Gamma}_s (\boldsymbol{\Gamma}_s^\top \boldsymbol{\Gamma}_s + \alpha_{w2s} N \boldsymbol{\Lambda}_s^{-1})^{-1} (\boldsymbol{\Gamma}_s^\top \boldsymbol{\Gamma}_s + \alpha_{w2s} N \boldsymbol{\Lambda}_s^{-1})^{-1} \boldsymbol{\Gamma}_s^\top \boldsymbol{\Gamma}_w \right) \mathbb{E}_{\tilde{\mathcal{S}}} [\boldsymbol{\zeta} \boldsymbol{\zeta}^\top] \right) \\ &= \text{tr} \left(\mathbb{E}_{\mathcal{S}_s} \left(\mathbf{V}_w^\top \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} + \alpha_{w2s} N \boldsymbol{\Sigma}_s^{-1})^{-2} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \mathbf{V}_w \right) \mathbb{E}_{\tilde{\mathcal{S}}} [\boldsymbol{\zeta} \boldsymbol{\zeta}^\top] \right) \\ &= \text{tr} \left(\mathbb{E}_{\mathcal{S}_s} \left(\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} + \alpha_{w2s} N \boldsymbol{\Sigma}_s^{-1})^{-2} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \right) \mathbb{E}_{\tilde{\mathcal{S}}} [\mathbf{V}_w \boldsymbol{\zeta} \boldsymbol{\zeta}^\top \mathbf{V}_w^\top] \right). \end{aligned}$$

Notice that $(\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} + \alpha_{w2s} N \boldsymbol{\Sigma}_s^{-1})^2 \succeq \alpha_{w2s} N (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \boldsymbol{\Sigma}_s^{-1} + \boldsymbol{\Sigma}_s^{-1} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma})$. Since matrix inversion is convex, a Jensen-type inequality implies that

$$\begin{aligned} &\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} + \alpha_{w2s} N \boldsymbol{\Sigma}_s^{-1})^{-2} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \\ &\preceq \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} (\alpha_{w2s} N (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \boldsymbol{\Sigma}_s^{-1} + \boldsymbol{\Sigma}_s^{-1} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma}))^\dagger \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \\ &= \frac{1}{2\alpha_{w2s} N} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \left(\frac{1}{2} (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \boldsymbol{\Sigma}_s^{-1} + \boldsymbol{\Sigma}_s^{-1} \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma}) \right)^\dagger \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \\ &\preceq \frac{1}{4\alpha_{w2s} N} (\boldsymbol{\Gamma}^\top \boldsymbol{\Gamma} \boldsymbol{\Sigma}_s + \boldsymbol{\Sigma}_s \boldsymbol{\Gamma}^\top \boldsymbol{\Gamma}). \end{aligned}$$

Therefore,

$$\mathbb{E}_{\mathcal{S}_s} \left(\mathbf{\Gamma}^\top \mathbf{\Gamma} (\mathbf{\Gamma}^\top \mathbf{\Gamma} + \alpha_{w2s} N \mathbf{\Sigma}_s^{-1})^{-2} \mathbf{\Gamma}^\top \mathbf{\Gamma} \right) \preceq \frac{1}{4\alpha_{w2s} N} \mathbb{E}_{\mathcal{S}_s} [\mathbf{\Gamma}^\top \mathbf{\Gamma} \mathbf{\Sigma}_s + \mathbf{\Sigma}_s \mathbf{\Gamma}^\top \mathbf{\Gamma}] = \frac{1}{2\alpha_{w2s} N} \mathbf{\Sigma}_s.$$

Meanwhile, we observe that

$$\begin{aligned} \mathbb{E}_{\tilde{\mathcal{S}}} [\mathbf{V}_w \boldsymbol{\zeta} \boldsymbol{\zeta}^\top \mathbf{V}_w^\top] &= \mathbb{E}_{\tilde{\mathcal{S}}} \left[\boldsymbol{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{z}} \tilde{\mathbf{z}}^\top \tilde{\mathbf{\Phi}}_w \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \boldsymbol{\Sigma}_w^{1/2} \right] \\ &= \sigma^2 \mathbb{E}_{\tilde{\mathcal{S}}} \left[\boldsymbol{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \boldsymbol{\Sigma}_w^{1/2} \right], \end{aligned}$$

where

$$\left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \preceq \frac{1}{2\alpha_w n} \mathbf{I}_d.$$

Therefore, we have

$$\mathbb{E}_{\tilde{\mathcal{S}}} [\mathbf{V}_w \boldsymbol{\zeta} \boldsymbol{\zeta}^\top \mathbf{V}_w^\top] \preceq \sigma^2 \mathbb{E}_{\tilde{\mathcal{S}}} \left[\boldsymbol{\Sigma}_w^{1/2} \left(\frac{1}{2\alpha_w n} \mathbf{I}_d \right) \boldsymbol{\Sigma}_w^{1/2} \right] = \frac{\sigma^2}{2\alpha_w n} \mathbf{\Sigma}_w.$$

Overall, the variance of f_{w2s} can be upper bounded as

$$\begin{aligned} \text{Var}(f_{w2s}) &= \text{tr} \left(\mathbb{E}_{\mathcal{S}_s} \left(\mathbf{\Gamma}^\top \mathbf{\Gamma} (\mathbf{\Gamma}^\top \mathbf{\Gamma} + \alpha_{w2s} N \mathbf{\Sigma}_s^{-1})^{-2} \mathbf{\Gamma}^\top \mathbf{\Gamma} \right) \mathbb{E}_{\tilde{\mathcal{S}}} [\mathbf{V}_w \boldsymbol{\zeta} \boldsymbol{\zeta}^\top \mathbf{V}_w^\top] \right) \\ &\leq \frac{\sigma^2 \text{tr}(\mathbf{\Sigma}_s \mathbf{\Sigma}_w)}{4(\alpha_w n)(\alpha_{w2s} N)}. \end{aligned} \quad (20)$$

Bias. Let $\boldsymbol{\xi} = \boldsymbol{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{f}}_* \in \mathbb{R}^d$, whose randomness comes from $\tilde{\mathcal{S}}$ only, independent of \mathcal{S}_x . Recall from (19), the bias term (19) can be decomposed as

$$\begin{aligned} \text{Bias}(f_{w2s}) &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Phi}_s \left(\mathbf{\Phi}_s^\top \mathbf{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d \right)^{-1} \mathbf{\Phi}_s^\top \mathbf{\Phi}_w \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{f}}_* - \mathbf{f}_* \right\|_2^2 \right] \\ &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left(\left\| \mathbf{\Phi}_s \left(\mathbf{\Phi}_s^\top \mathbf{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d \right)^{-1} \mathbf{\Phi}_s^\top \mathbf{\Gamma} \boldsymbol{\xi} - \mathbf{\Phi}_s \mathbf{\Phi}_s^\dagger \mathbf{f}_* \right\|_2^2 + \left\| (\mathbf{I}_N - \mathbf{\Phi}_s \mathbf{\Phi}_s^\dagger) \mathbf{f}_* \right\|_2^2 \right) \right], \end{aligned}$$

where by Lemma B.1 and (14)

$$\mathbb{E}_{\mathcal{S}_x} \left[\frac{1}{N} \left\| (\mathbf{I}_N - \mathbf{\Phi}_s \mathbf{\Phi}_s^\dagger) \mathbf{f}_* \right\|_2^2 \right] = \frac{\rho_s(N)}{N} \leq \rho_s = 0.$$

Therefore, with $\boldsymbol{\xi} = \boldsymbol{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{f}}_*$, we have

$$\text{Bias}(f_{w2s}) = \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Phi}_s \left(\mathbf{\Phi}_s^\top \mathbf{\Phi}_s + \alpha_{w2s} N \mathbf{I}_d \right)^{-1} \mathbf{\Phi}_s^\top \mathbf{\Gamma} \boldsymbol{\xi} - \mathbf{\Phi}_s \mathbf{\Phi}_s^\dagger \mathbf{f}_* \right\|_2^2 \right].$$

Recall that $\mathbf{f}_* = \mathbf{\Gamma} \boldsymbol{\Sigma}_*^{1/2} \boldsymbol{\theta}_*$ and $\mathbf{\Phi}_s = \mathbf{\Gamma} \boldsymbol{\Sigma}_s^{1/2} = \mathbf{\Gamma}_s \boldsymbol{\Lambda}_s^{1/2} \mathbf{V}_s^\top$. Then, we can express the bias term as

$$\begin{aligned} \text{Bias}(f_{w2s}) &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Gamma} \left(\mathbf{\Gamma}^\top \mathbf{\Gamma} + \alpha_{w2s} N \mathbf{\Sigma}_s^{-1} \right)^{-1} \mathbf{\Gamma}^\top \mathbf{\Gamma} \boldsymbol{\xi} - \mathbf{\Gamma} \mathbf{\Gamma}^\dagger \mathbf{f}_* \right\|_2^2 \right] \\ &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Gamma} \boldsymbol{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \mathbf{\Gamma} \left(\mathbf{\Gamma}^\top \mathbf{\Gamma} + \alpha_{w2s} N \mathbf{\Sigma}_s^{-1} \right)^{-1} \mathbf{\Gamma}^\top \mathbf{\Gamma} \boldsymbol{\xi} \right\|_2^2 \right] \\ &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Gamma} \left(\boldsymbol{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \boldsymbol{\xi} \right) + \mathbf{\Gamma} \left(\mathbf{I}_d - \left(\mathbf{\Gamma}^\top \mathbf{\Gamma} + \alpha_{w2s} N \mathbf{\Sigma}_s^{-1} \right)^{-1} \mathbf{\Gamma}^\top \mathbf{\Gamma} \right) \boldsymbol{\xi} \right\|_2^2 \right] \end{aligned}$$

By the Woodbury matrix identity, we have

$$\mathbf{I}_d - (\mathbf{\Gamma}^\top \mathbf{\Gamma} + \alpha_{w2s} N \mathbf{\Sigma}_s^{-1})^{-1} \mathbf{\Gamma}^\top \mathbf{\Gamma} = \left(\mathbf{I}_d + \frac{1}{\alpha_{w2s} N} \mathbf{\Sigma}_s \mathbf{\Gamma}^\top \mathbf{\Gamma} \right)^{-1}. \quad (21)$$

Therefore, we have

$$\mathbf{Bias}(f_{w2s}) = \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\underbrace{\frac{1}{N} \left\| \mathbf{\Gamma} \left(\mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \boldsymbol{\xi} \right) \right\|_2^2}_{\text{Term A}} + \underbrace{\mathbf{\Gamma} \left(\mathbf{I}_d + \frac{1}{\alpha_{w2s} N} \mathbf{\Sigma}_s \mathbf{\Gamma}^\top \mathbf{\Gamma} \right)^{-1} \boldsymbol{\xi} \left\|_2^2}_{\text{Term B}} \right]. \quad (22)$$

For Term A, notice that $\boldsymbol{\xi} = \mathbf{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{f}}_*$ implies

$$\begin{aligned} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \boldsymbol{\xi} &= \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \mathbf{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{f}}_* \\ &= \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \left(\tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} + \alpha_w n \mathbf{\Sigma}_w^{-1} \right)^{-1} \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* \\ &= \left(\mathbf{I}_d - \left(\tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} + \alpha_w n \mathbf{\Sigma}_w^{-1} \right)^{-1} \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \right) \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* \\ &= \left(\mathbf{I}_d + \frac{1}{\alpha_w n} \mathbf{\Sigma}_w \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \right)^{-1} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_*, \end{aligned}$$

where the last equality follows from Woodbury matrix identity as in (21). Therefore,

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Gamma} \left(\mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \boldsymbol{\xi} \right) \right\|_2^2 \right] &= \mathbb{E}_{\tilde{\mathcal{S}}} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Gamma}} \left(\mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \boldsymbol{\xi} \right) \right\|_2^2 \right] \\ &= \mathbb{E}_{\tilde{\mathcal{S}}} \left[\frac{1}{n} \left\| \tilde{\mathbf{\Gamma}} \left(\mathbf{I}_d + \frac{1}{\alpha_w n} \mathbf{\Sigma}_w \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \right)^{-1} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* \right\|_2^2 \right]. \end{aligned}$$

Since

$$\left(\mathbf{I}_d + \frac{1}{\alpha_w n} \mathbf{\Sigma}_w \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \right)^{-1} \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \left(\mathbf{I}_d + \frac{1}{\alpha_w n} \mathbf{\Sigma}_w \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \right)^{-1} \preceq \frac{\alpha_w n}{2} \mathbf{\Sigma}_w^{-1},$$

we have

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Gamma} \left(\mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* - \boldsymbol{\xi} \right) \right\|_2^2 \right] &\leq \frac{1}{n} \text{tr} \left(\frac{\alpha_w n}{2} \mathbf{\Sigma}_w^{-1} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* \boldsymbol{\theta}_*^\top \mathbf{\Sigma}_*^{1/2} \right) \\ &= \frac{\alpha_w}{2} \left\| \mathbf{\Sigma}_w^{-1/2} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* \right\|_2^2. \end{aligned} \quad (23)$$

For Term B, leveraging Woodbury matrix identity as in (21), we notice that

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \mathbf{\Gamma} \left(\mathbf{I}_d + \frac{1}{\alpha_{w2s} N} \mathbf{\Sigma}_s \mathbf{\Gamma}^\top \mathbf{\Gamma} \right)^{-1} \boldsymbol{\xi} \right\|_2^2 \right] &\leq \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \text{tr} \left(\frac{\alpha_{w2s} N}{2} \mathbf{\Sigma}_s^{-1} \boldsymbol{\xi} \boldsymbol{\xi}^\top \right) \right] \\ &= \frac{\alpha_{w2s}}{2} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \mathbf{\Sigma}_s^{-1/2} \mathbf{\Sigma}_w^{1/2} \left(\tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{\Phi}}_w + \alpha_w n \mathbf{I}_d \right)^{-1} \tilde{\mathbf{\Phi}}_w^\top \tilde{\mathbf{f}}_* \right\|_2^2 \right] \\ &= \frac{\alpha_{w2s}}{2} \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\left\| \mathbf{\Sigma}_s^{-1/2} \left(\tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} + \alpha_w n \mathbf{\Sigma}_w^{-1} \right)^{-1} \tilde{\mathbf{\Gamma}}^\top \tilde{\mathbf{\Gamma}} \mathbf{\Sigma}_*^{1/2} \boldsymbol{\theta}_* \right\|_2^2 \right] \end{aligned}$$

Since $(\tilde{\Gamma}^\top \tilde{\Gamma} + \alpha_w n \Sigma_w^{-1})^{-1} \tilde{\Gamma}^\top \tilde{\Gamma} \preceq \mathbf{I}_d$, we know that

$$\mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \Gamma \left(\mathbf{I}_d + \frac{1}{\alpha_w 2s N} \Sigma_s \Gamma^\top \Gamma \right)^{-1} \xi \right\|_2^2 \right] \leq \frac{\alpha_w 2s}{2} \left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2. \quad (24)$$

Combining (22), (23), and (24), we can upper bound the bias term as

$$\begin{aligned} \text{Bias}(f_{w2s}) &= \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \underbrace{\Gamma (\Sigma_*^{1/2} \theta_* - \xi)}_{\text{Term A}} + \underbrace{\Gamma \left(\mathbf{I}_d + \frac{1}{\alpha_w 2s N} \Sigma_s \Gamma^\top \Gamma \right)^{-1} \xi}_{\text{Term B}} \right\|_2^2 \right] \\ &\leq 2 \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \Gamma (\Sigma_*^{1/2} \theta_* - \xi) \right\|_2^2 \right] + 2 \mathbb{E}_{\mathcal{S}_x, \tilde{\mathcal{S}}} \left[\frac{1}{N} \left\| \Gamma \left(\mathbf{I}_d + \frac{1}{\alpha_w 2s N} \Sigma_s \Gamma^\top \Gamma \right)^{-1} \xi \right\|_2^2 \right] \\ &\leq \alpha_w \left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2 + \alpha_w 2s \left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2. \end{aligned} \quad (25)$$

Variance-bias tradeoff. Overall, by (20) and (25), we have

$$\begin{aligned} \text{Var}(f_{w2s}) &\leq \frac{\sigma^2 \text{tr}(\Sigma_s \Sigma_w)}{4(\alpha_w n)(\alpha_w 2s N)}, \\ \text{Bias}(f_{w2s}) &\leq \alpha_w \left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2 + \alpha_w 2s \left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2. \end{aligned}$$

The upper bound the excess risk $\text{ER}(f_{w2s}) = \text{Var}(f_{w2s}) + \text{Bias}(f_{w2s})$ is minimized by taking

$$\alpha_w = \left(\frac{\sigma^2 \text{tr}(\Sigma_s \Sigma_w)}{4nN} \frac{\left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2}{\left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^4} \right)^{1/3}, \quad \alpha_{w2s} = \left(\frac{\sigma^2 \text{tr}(\Sigma_s \Sigma_w)}{4nN} \frac{\left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2}{\left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^4} \right)^{1/3},$$

which leads to the optimal upper bound for the excess risk:

$$\text{ER}(f_{w2s}) \leq 3 \left(\frac{\sigma^2 \text{tr}(\Sigma_s \Sigma_w)}{4nN} \left\| \Sigma_s^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2 \left\| \Sigma_w^{-1/2} \Sigma_*^{1/2} \theta_* \right\|_2^2 \right)^{1/3}.$$

□

D. Canonical angles

In this section, we review the concept of canonical angles between two subspaces that connect the formal definition of the correlation dimension $d_{s \wedge w} = \|\mathbf{V}_s^\top \mathbf{V}_w\|_F^2$ in Definition 2.4 to the intuitive notion of the alignment between the corresponding subspaces \mathcal{V}_s and \mathcal{V}_w in the introduction: $\sum \cos(\angle(\mathcal{V}_s, \mathcal{V}_w)) = \|\mathbf{V}_s^\top \mathbf{V}_w\|_F^2$.

Definition D.1 (Canonical angles (Golub & Van Loan, 2013), adapting from (Dong et al., 2024a)). Let $\mathcal{V}_s, \mathcal{V}_w \subseteq \mathbb{R}^d$ be two subspaces with dimensions $\dim(\mathcal{V}_s) = d_s$ and $\dim(\mathcal{V}_w) = d_w$ (assuming $d_w \geq d_s$ without loss of generality). The canonical angles $\angle(\mathcal{V}_s, \mathcal{V}_w) = \text{diag}(\nu_1, \dots, \nu_{d_s})$ are d_s angles that jointly measure the alignment between \mathcal{V}_s and \mathcal{V}_w , defined recursively as follows:

$$\begin{aligned} \mathbf{u}_i, \mathbf{v}_i &\triangleq \arg\max \mathbf{u}_i^* \mathbf{v}_i \\ \text{s.t. } \mathbf{u}_i &\in \left(\mathcal{V}_s \setminus \text{span}\{\mathbf{u}_\ell\}_{\ell=1}^{i-1} \right) \cap \mathbb{S}^{d-1}, \\ \mathbf{v}_i &\in \left(\mathcal{V}_w \setminus \text{span}\{\mathbf{v}_\ell\}_{\ell=1}^{i-1} \right) \cap \mathbb{S}^{d-1} \\ \cos(\nu_i) &= \mathbf{u}_i^* \mathbf{v}_i \quad \forall i = 1, \dots, k, \end{aligned}$$

such that $0 \leq \nu_1 \leq \dots \leq \nu_k \leq \pi/2$.

Given two subspaces $\mathcal{V}_s, \mathcal{V}_w \subseteq \mathbb{R}^d$, let $\mathbf{V}_s \in \mathbb{R}^{d \times d_s}$ and $\mathbf{V}_w \in \mathbb{R}^{d \times d_w}$ be the matrices whose columns form orthonormal bases for \mathcal{V}_s and \mathcal{V}_w , respectively. Then, the canonical angles $\angle(\mathcal{V}_s, \mathcal{V}_w)$ are determined by the singular values of $\mathbf{V}_s^\top \mathbf{V}_w$ (Björck & Golub, 1973, §3):

$$\cos(\angle_i(\mathcal{V}_s, \mathcal{V}_w)) = \sigma_i(\mathbf{V}_s^\top \mathbf{V}_w) \quad \forall i = 1, \dots, d_s,$$

where $\sigma_i(\mathbf{V}_s^\top \mathbf{V}_w)$ denotes the i -th singular value of $\mathbf{V}_s^\top \mathbf{V}_w$.

In particular, since $\mathbf{V}_s, \mathbf{V}_w$ consist of orthonormal columns, the singular values of $\mathbf{V}_s^\top \mathbf{V}_w$ fall in $[0, 1]$, and therefore,

$$d_{s \wedge w} = \sum \cos(\angle(\mathcal{V}_s, \mathcal{V}_w)) = \|\mathbf{V}_s^\top \mathbf{V}_w\|_F^2 \in [0, \min\{d_s, d_w\}].$$

E. Additional experiments

E.1. Additional experiments and details on UTKFace regression

This section provides some additional details and results for the UTKFace regression experiments in Section 4.2.




Figure 7. Scaling for MSE on UTKFace with CLIP-B32 as the strong student and ResNet18 as the weak teacher

We summarize the relevant dimensionality in Table 1. We observe the following:

- The intrinsic dimensions of the pretrained features are significantly smaller than the ambient feature dimensions, which is consistent with our theoretical analysis and the empirical observations in (Aghajanyan et al., 2021).
- The correlation dimensions $d_{s \wedge w}$ are considerably smaller than the corresponding intrinsic dimensions, indicating that the subspaces spanned by the weak and strong features are not aligned in practice. As shown in Section 4.2, such discrepancies in the weak and strong features facilitate W2S generalization.

For reference, we provide the scaling for MSE losses of three representative teacher-student pairs in Figures 7 to 9.

- It is worth highlighting that while the MSE loss of f_{w2s} monotonically decreases with respect to both sample sizes n, N , the different rates of convergence compared to f_w, f_s , and f_c lead to the distinct scaling behavior of the relative W2S performance (PGR and OPR) with respect to n versus N in Figures 5 and 6.
- When the strong student has a lower intrinsic dimension than the weak teacher (cf. Figure 7 versus Figures 8 and 9), $d_s < d_w$, the W2S model f_{w2s} tends to achieve better generalization in terms of the test MSE. This is consistent with our analysis in Section 3.1.
- When $d_s < d_w$, the W2S model f_{w2s} tends to achieve (slightly) better generalization for (slightly) smaller correlation dimension $d_{s \wedge w}$ (cf. Figure 8 versus Figure 9), again coinciding with our analysis in Section 3.1.




Figure 8. Scaling for MSE on UTKFace with CLIP-B32 as the strong student and ResNet50 as the weak teacher




Figure 9. Scaling for MSE on UTKFace with CLIP-B32 as the strong student and ResNet152 as the weak teacher

 Table 1. Summary of the pretrained feature dimensions, along with the intrinsic dimensions d_s , d_w and correlation dimensions $d_{s \wedge w}$ (with respect to the strong student CLIP-B32) computed over the entire UTKFace dataset (including training and testing).

Pretrained Model	Feature Dimension	Intrinsic Dimension ($\tau = 0.01$)	Correlation Dimension
ResNet18	512	194	167.64
ResNet34	512	150	129.97
ResNet50	2048	522	301.06
ResNet101	2048	615	354.52
ResNet152	2048	589	339.90
CLIP-B32	768	443	×

- W2S generalization generally happens (*i.e.* f_{w2s} is able to outperform f_w) with sufficiently large sample sizes n, N . However, as the labeled sample size n increases, the test MSE of f_{w2s} converges slower than that of the strong baseline and ceiling models, f_s and f_c , leading to the inverse scaling for PGR and OPR with respect to n in Figures 5 and 6. When n is too large, the W2S model f_{w2s} may not be able to achieve better generalization than the strong baseline f_s .

E.2. Experiments on image classification

Dataset. ColoredMNIST (Arjovsky et al., 2019) consists of groups of different colors and reassign the label to be binary (digits 0-4 vs 5-9). We pool together the groups to form one dataset. The choice is to bring diversity to the feature space with additional color features and thus potential feature discrepancies. We hold out a test set of 7000 samples and use the rest 63000 samples for training.

Linear probing over pretrained features. We fix a strong student as DINOv2-s14 (Oquab et al., 2024) and vary the weak teacher among the ResNet-d series and ResNet series (ResNet18D, ResNet34D, ResNet101, ResNet152) (He et al., 2019; 2016). We replace ResNet18 and ResNet34 used in Section 4.2 to experiment on weak models with similar intrinsic dimensions but different correlation dimensions. We treat the backbone of the models (excluding the classification layer) as ϕ_s and ϕ_w and finetune them via linear probing. We train the models with cross-entropy loss and AdamW optimizer. We tune the training hyperparameters of weak and strong models using a validation set and train them for 800 steps with a learning rate 1e-3 and weight decay 1e-6.

Table 2. Summary of the pretrained feature dimensions, along with the intrinsic dimensions d_s, d_w and correlation dimensions $d_{s \wedge w}$ (with respect to the strong student DINOv2-S14) computed over the entire ColoredMNIST dataset (including training and testing).

Pretrained Model	Feature Dimension	Intrinsic Dimension ($\tau = 0.01$)	Correlation Dimension
ResNet-18-D	512	117	6.23
ResNet-34-D	512	127	7.07
ResNet101	2048	121	1.74
ResNet152	2048	128	1.88
DINOv2-S14	384	28	\times

$d_w = 117$ (ResNet18D), 127 (ResNet34D), 121 (ResNet101), 128 (ResNet152), $d_s = 28$ (DINOv2-S14)

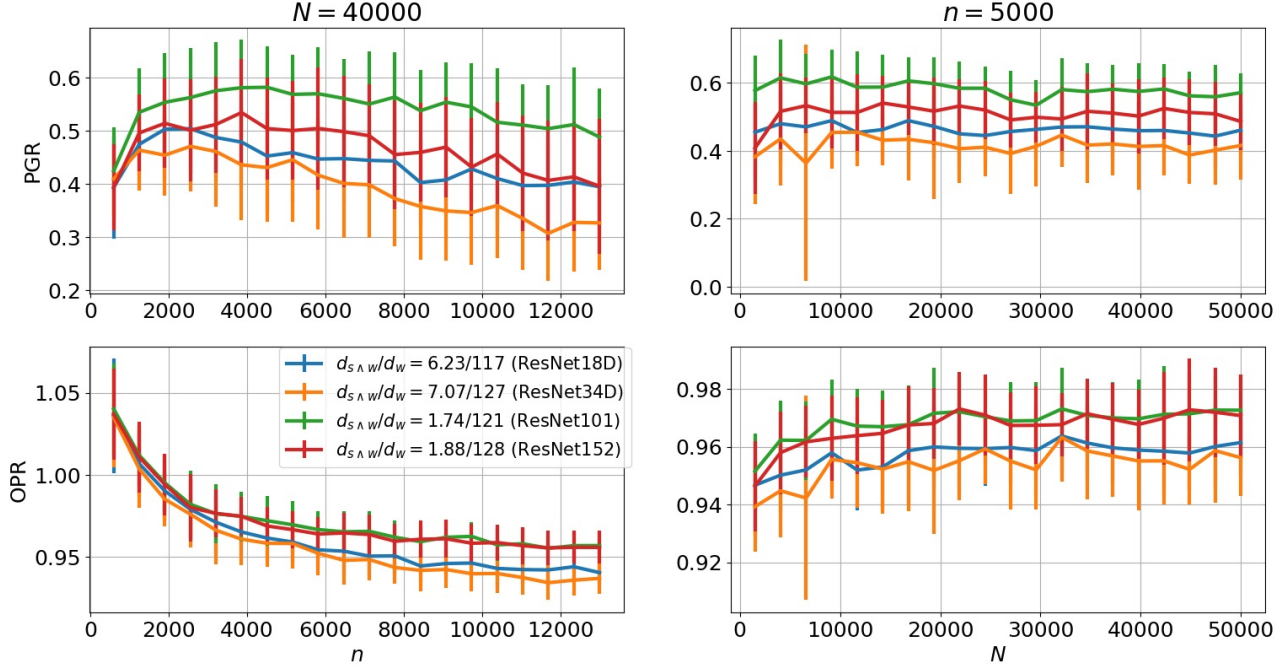


Figure 10. Scaling for PGR and OPR of different weak teachers with a fixed strong student on ColoredMNIST.

Discrepancies lead to better W2S. Figure 10 shows the scaling of PGR and OPR with respect to the sample sizes n, N for different weak teachers in the ResNet series with respect to a fixed student, CLIP-B32. As in Section 4.2, we

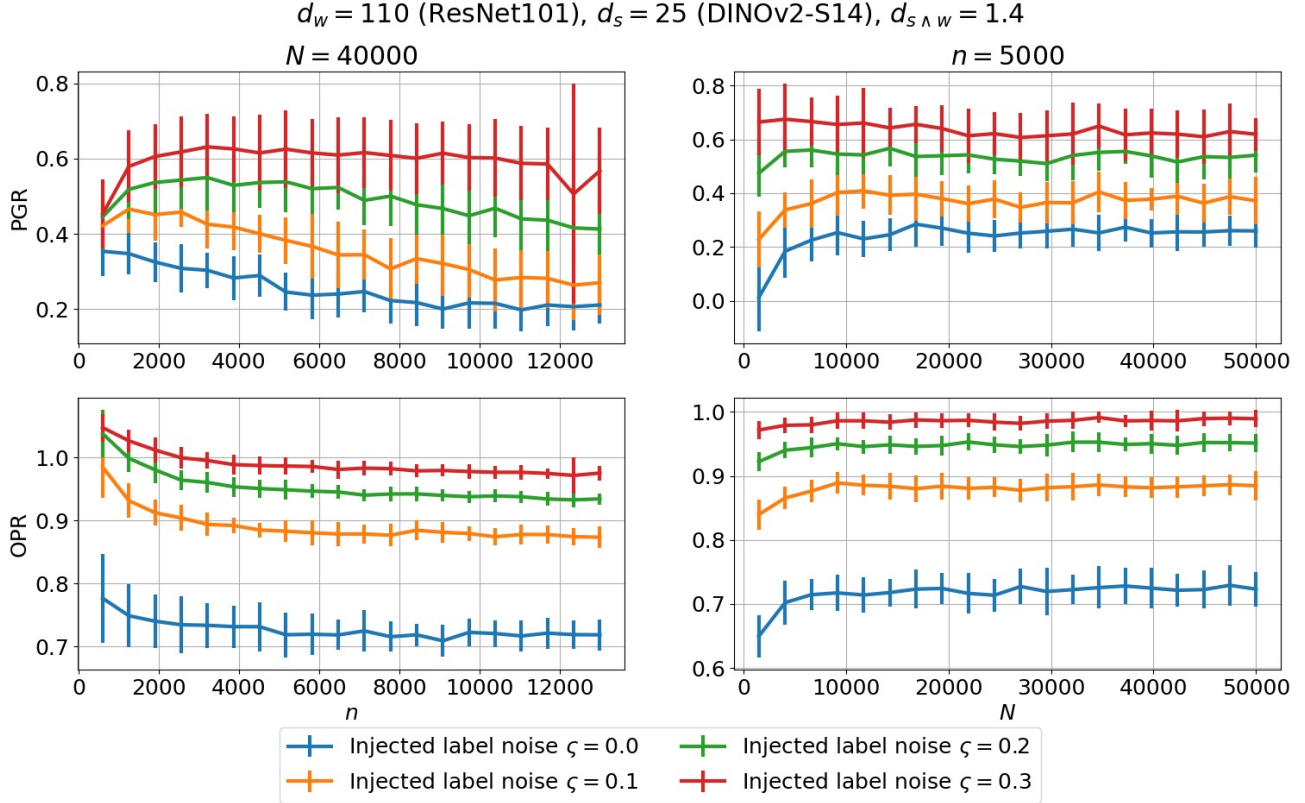


Figure 11. Scaling for **PGR** and **OPR** of W2S on ColoredMNIST with injected label noise.

observe that with similar intrinsic dimensions d_s, d_w , W2S achieves better relative performance in terms of **PGR** and **OPR** when the correlation dimension $d_{s \wedge w}$ is smaller.

Variance reduction is a key advantage of W2S. We inject noise to the labels of the original ColoredMNIST training samples by randomly flipping the ground truth labels with probability $\zeta \in [0, 1]$ (following (Arjovsky et al., 2019)). Figure 11 shows the scaling of **PGR** and **OPR** with respect to n and N when taking DINOv2-S14 as the strong student and ResNet101 as the weak teacher. We observe that the larger artificial label noise ζ leads to higher **PGR** and **OPR**.

E.3. Experiments on sentiment classification

Dataset. The Stanford Sentiment Treebank (Socher et al., 2013) is a corpus with fully labeled parse trees that allows for a complete analysis of the compositional effects of sentiment in language. The corpus is based on the dataset introduced by Pang & Lee (2005) and consists of 11,855 single sentences extracted from movie reviews. It was parsed with the Stanford parser and includes a total of 215,154 unique phrases from those parse trees, each annotated by 3 human judges. We conduct binary classification experiments on full sentences (negative or somewhat negative versus somewhat positive or positive, with neutral sentences discarded), the so-called SST-2 dataset, and split the dataset into training and testing sets of sizes 63000 and 4349. Generalization errors are estimated with the 0-1 loss over the test set.

Full finetuning. We fix the strong student as Electra-base-discriminator (Clark et al., 2020) and vary the weak teacher among the Bert series (Turc et al., 2019) (Bert-Tiny, Bert-Mini, Bert-Small, Bert-Medium). With manageable model sizes, we conduct full finetuning experiments following the setup in (Burns et al., 2024). We use the standard cross-entropy loss for supervised finetuning. When training strong students on weak labels (W2S), we use the confidence-weighted loss proposed by (Burns et al., 2024), which is suggested to be able to improve weak-to-strong generalization on many NLP tasks. All training is conducted via Adam optimizers (Kingma & Ba, 2014) with a learning rate of $5e-5$, a cosine learning rate schedule, and 40 warmup steps. We train for 3 epochs, which is sufficient for the train and validation losses to stabilize.

Intrinsic dimension. The intrinsic dimensions d_w, d_s are measured based on the Structure-Aware Intrinsic Dimension (SAID) method proposed by (Aghajanyan et al., 2021). We first train the full models on the whole training set, and then train the models with only d trainable parameters based on SAID transformation. The d_w or d_s are the smallest number of parameters under SAID that is necessary to retain 90% accuracy of the full models. Here, the 90% accuracy is a common threshold used to estimate intrinsic dimensions in the literature (Li et al., 2018).




Figure 12. Scaling for PGR and OPR of different weak teachers with a fixed strong student on SST-2.

Correlation Dimension. Let $D_s, D_w \in \mathbb{N}$ be the finetunable parameter counts of the strong and weak models, respectively. For full FT whose dynamics fall in the kernel regime, as explained in Remark 2.5, the strong and weak “features” become the gradients⁶, $\Phi_s = \nabla_{\theta} f_s(\mathbf{X} | \theta_s^{(0)}) \in \mathbb{R}^{N \times D_s}$ and $\Phi_w = \nabla_{\theta} f_w(\mathbf{X} | \theta_w^{(0)}) \in \mathbb{R}^{N \times D_w}$, of the respective models at the pretrained initialization, $\theta_s^{(0)} \in \mathbb{R}^{D_s}$ and $\theta_w^{(0)} \in \mathbb{R}^{D_w}$.

A practical challenge is that D_s, D_w, N are all huge for full FT on most NLP tasks, making it infeasible to compute the $D_s \times D_s$ and $D_w \times D_w$ Gram matrices and their spectral decompositions. As a remedy, we leverage the significantly lower intrinsic dimensions $d_s \ll D_s, d_w \ll D_w$ (see Table 2) to accelerate estimation of $d_{s \wedge w}$ via sketching (Halko et al., 2011; Woodruff et al., 2014).

- (i) We first reduce both D_s, D_w to the same lower dimension $D = 0.01 \min\{D_s, D_w\}$ (with $D \gg \max\{d_s, d_w\}$) by subsampling columns of Φ_s, Φ_w (uniformly for efficiency, or adaptively via sketching-based interpolative decomposition (Dong & Martinsson, 2023) when affordable) to obtain $\Phi'_s, \Phi'_w \in \mathbb{R}^{N \times D}$.
- (ii) Then, we use randomized rangefinder (Halko et al., 2011, Algorithm 4.1) to approximate the first d_s, d_w right singular vectors, $\mathbf{V}_s \in \mathbb{R}^{D \times d_s}$ and $\mathbf{V}_w \in \mathbb{R}^{D \times d_w}$, of Φ'_s, Φ'_w . Taking the evaluation of \mathbf{V}_s as an example, we draw a Gaussian random matrix $\mathbf{G}_s \in \mathbb{R}^{d_s \times D}$ and compute the orthonormalization $\mathbf{V}_s = \text{ortho}(\Phi_s'^T \mathbf{G}_s)$ via QR decomposition.

⁶Notice that f_s, f_w are scalar-valued functions for binary classification tasks like SST-2, and thus the gradients $\nabla_{\theta} f_s$ and $\nabla_{\theta} f_w$ are row vectors. For multi-class classification tasks where f_s, f_w output vectors of logits, a common heuristic to keep Φ_s, Φ_w as matrices of manageable sizes (in contrast to tensors) is to replace gradients of the models, $\nabla_{\theta} f_s$ and $\nabla_{\theta} f_w$, with gradients of MSE losses at the pretrained initialization. The gradients of MSE can be viewed as a weighted sum of the model gradients for each class.

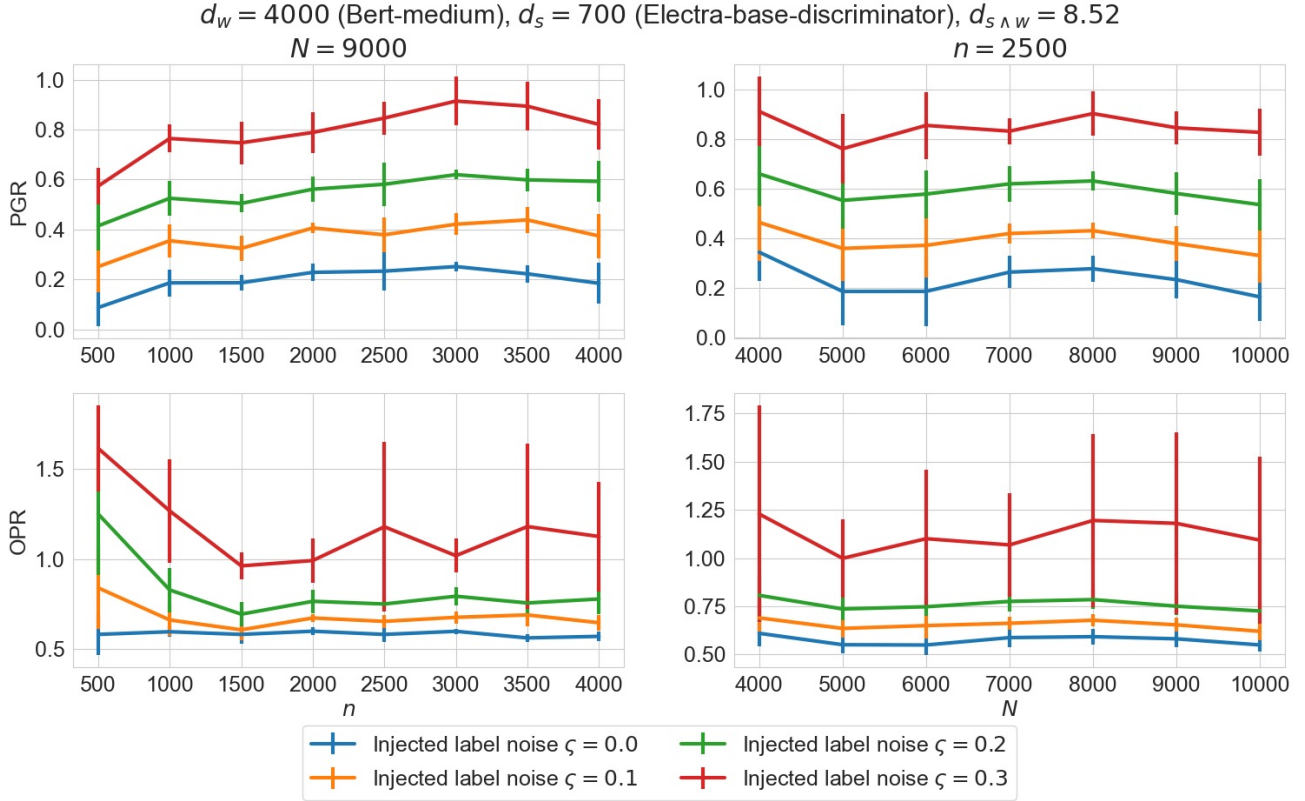


Figure 13. Scaling for **PGR** and **OPR** of W2S on SST-2 with injected label noise.

(iii) Finally, we compute the correlation dimension $d_{s \wedge w} = \|\mathbf{V}_s^\top \mathbf{V}_w\|_F^2$.

Table 3. Summary of finetunable parameter counts D_s, D_w , intrinsic dimensions d_s, d_w , and correlation dimensions $d_{s \wedge w}$ (with respect to the strong student **Electra**) computed over the entire SST-2 dataset (including training and testing).

Pretrained Model	D_s, D_w	Intrinsic Dimension ($\tau = 0.01$)	Correlation Dimension
Bert-Tiny	4.4M	7000	81.13
Bert-Mini	11.2M	8500	38.67
Bert-Small	28.8M	8000	26.19
Bert-Medium	41.4M	4000	8.52
Electra	109.5M	700	×

Discrepancies lead to better W2S. Figure 12 shows the scaling of **PGR** and **OPR** with respect to n and N for different $d_{s \wedge w}$. As in Section 4.2 and Appendix E.2, we observe the better relative W2S performance in terms of **PGR** and **OPR** when $d_{s \wedge w}/d_w$ is smaller.

Variance reduction is a key advantage of W2S. We inject noise to the labels of training samples by randomly flipping labels with probability $\zeta = 0, 0.1, 0.2, 0.3$. Figure 13 shows the scaling of **PGR** and **OPR** with respect to n and N when taking **Electra** as the strong student and **Bert-Medium** as the weak teacher. We observe that the larger artificial label noise ζ leads to higher **PGR** and **OPR**.