

OPTIMAL CRITERION FOR FEATURE LEARNING OF TWO-LAYER LINEAR NEURAL NETWORK IN HIGH DIMENSIONAL INTERPOLATION REGIME

Keita Suzuki

University of Tokyo and Preferred Networks, Inc.
keitasuzuki@preferred.jp

Taiji Suzuki

University of Tokyo and RIKEN AIP
taiji@mist.i.u-tokyo.ac.jp

ABSTRACT

Deep neural networks with feature learning have shown surprising generalization performance in high dimensional settings, but it has not been fully understood how and when they enjoy the benefit of feature learning. In this paper, we theoretically analyze the statistical properties of the benefit from feature learning in a two-layer linear neural network with multiple outputs in a high-dimensional setting. For that purpose, we propose a new criterion that allows feature learning of a two-layer linear neural network in a high-dimensional setting. Interestingly, we can show that models with smaller values of the criterion generalize even in situations where normal ridge regression fails to generalize. This is because the proposed criterion contains a proper regularization for the feature mapping and acts as an upper bound on the predictive risk. As an important characterization of the criterion, the two-layer linear neural network that minimizes this criterion can achieve the optimal Bayes risk that is determined by the distribution of the true signals across the multiple outputs. To the best of our knowledge, this is the first study to specifically identify the conditions under which a model obtained by proper feature learning can outperform normal ridge regression in a high-dimensional multiple-output linear regression problem.

1 INTRODUCTION

Thanks to the benefit of feature learning, neural networks have demonstrated superior generalization performance across numerous high-dimensional problems, including image classification and natural language processing. This empirical observation is surprising as it conflicts with conventional statistical theory, which claims a model shouldn't be too complex to predict unseen randomly labeled data well. Therefore it has drawn great attention from the community of learning theory (Belkin et al., 2018). To understand the generalization of high-dimensional statistical models, many researchers have studied the generalization mechanism of simpler models such as linear models, which are amenable to theoretical analysis in over-parametrized settings. Thanks to intensive research, we've gained an understanding of benign overfitting phenomenon in many simple models. In this phenomenon, models generalize well to unseen test data even when they interpolate training data with significant label noise, and we can see this in linear regression (Bartlett et al., 2020; Hastie et al., 2022; Muthukumar et al., 2020; Negrea et al., 2020; Xu and Hsu, 2019), ridge regression (Tsigler and Bartlett, 2020; Dobriban and Wager, 2018; Mei and Montanari, 2022), logistic regression (Montanari et al., 2019; Chatterji and Long, 2021), and kernel-based estimators (Liang and Rakhlin, 2020; Belkin et al., 2019; Li et al., 2022). However, in practical use, neural networks frequently surpass these simpler models in high-dimensional problem settings. Hence, understanding the benefits of feature learning and its dominance over simple models without feature learning remains crucial in order to elucidate the generalization of deep learning. For this reason, many researchers have studied feature learning of two-layer neural networks, which is the simplest case of neural networks. However, much of this research is confined to the setting of classification problems (Frei et al., 2022; Cao et al., 2022), with few studies dedicated to regression problems. This is because there is a difficulty in feature learning for over-parameterized regression problems that increasing the number of parameters excessively for feature mapping can lead to overfitting. Accordingly, it's crucial to apply regularization for parameters including feature mapping for appropriate

feature learning. As a regularization for such parameters, a considerable number of researchers are focusing on implicit regularization. This form of regularization is automatically induced by optimization methods and deep learning architectures. Several studies have examined the influence of implicit regularization on the generalization of two-layer neural networks in high-dimensional regression problems (Chatterji et al., 2022; Azulay et al., 2021). However, these analyses on implicit regularization aren't enough to fully understand the benefit of feature learning. This is because such implicit regularization does not always leverage the entire potential of feature learning, that is, it isn't necessarily an optimal regularization in terms of feature learning. In particular, the existing work on implicit regularization as listed above missed important properties of feature learning such as the Bayes risk optimality that will be shown in our study. Therefore, the goal of this work is to answer the following question:

Can we design an optimal regularization of feature mapping to fully exploit the benefit of feature learning and demonstrate its improvement over simpler models like ridge regression?

1.1 CONTRIBUTIONS

For the goal above, we consider a two-layer linear neural network which is the simplest setting of feature learning in a multi-output linear regression problem setting. We use the ridge estimator as the second-layer parameter and design a proper regularization on the first-layer parameter which enables feature learning. Our contribution can be summarized as follows:

- We propose a criterion that gives an upper bound of the predictive risk and show that the predictive risk for an estimator with smaller values of the criterion converges to zero. This criterion has a proper regularization term on the first-layer parameter.
- An important characterization of the proposed criterion is that the minimizer of the criterion achieves near-optimal Bayes risk under some conditions, which represents the adequacy of the criterion. This is a merit of considering a multi-output setting.
- We characterize when the normal ridge estimator does not generalize and show the two-layer neural network selected by the proposed criterion generalizes even in such situations using examples. Indeed, the predictive risk of such an estimator can converge to zero even when that of the naive ridge regression does not. This result demonstrates the benefit of feature learning and its dominance over simpler models in high-dimensional regression problems.

1.2 RELATED WORK

Benign overfitting of linear models. One line of research sought a theoretical understanding of the benign overfitting phenomenon in linear models. Hastie et al. (2022) gave an asymptotic evaluation of the bias and variance of linear models and random feature regression. Additionally, they demonstrated that the predictive risks associated with certain models exhibit a double descent curve as the ratio of parameter size to sample size becomes large. Bartlett et al. (2020) derived a tight upper bound for the predictive risk of the minimum-norm interpolator in an over-parametrized setting. They further elucidated the role of the input covariance matrix's spectrum in enabling the model to generalize well. Furthermore, Tsigler and Bartlett (2020) extended this result to ridge regression. Belkin et al. (2019) showed that the Nadaraya-Watson estimator with a singular kernel that perfectly fits training data enjoys the benign overfitting phenomenon, because thanks to singular kernels, this model fits the training data with spike-like shape. Xu and Hsu (2019) gave an evaluation of the predictive risk depending on the number of features in principal component regression and showed that the number of features shouldn't be reduced too much in high dimensional settings. Li et al. (2022) showed that in the framework of (Bartlett et al., 2020), noises added to features play a role of regularization and demonstrated that the double descent phenomenon occurs in this model. However, these methods are not optimal and can sometimes fail to generalize in our settings as this study demonstrates.

Benign overfitting of two-layer neural network. Many recent works have characterized the generalization performance of two-layer neural networks. Chatterji et al. (2022) studied the two-layer linear neural network setting we consider here and analyzed the generalization performance of a two-layer linear neural network trained by gradient flow. Frei et al. (2022) studied benign overfitting of a two-layer neural network with nonlinear activation function in classification problems. They showed that if the ground truth is linear, a two-layer neural network trained by gradient descent can achieve min-max optimal test error even if its training error is vanishingly small. Cao et al. (2022)

also studied benign overfitting in classification problem settings. They characterized the benign overfitting phenomenon in terms of signal-to-noise ratio in a convolutional neural network whose activation function is the power of ReLU. Additionally, they demonstrated that if the signal-to-noise ratio satisfies a certain condition, a two-layer CNN trained by gradient descent achieves arbitrary small training and test loss, and otherwise, the CNN can only achieve constant level test loss. Other works (Ba et al., 2022; Damian et al., 2022) showed that the two-layer neural networks with one gradient step on the first-layer parameter can outperform fixed kernel regression. Especially, Ba et al. (2022) highlighted the importance of a large learning rate in the early phase of training and demonstrated that if the learning rate is small, the two-layer neural network can't defeat the best linear model. However, these analyses handled only one-step gradient learning, as they were motivated by optimizability, and it's unknown what kind of statistical optimality the actual optimal solution satisfies in situations where $n \ll d$. Indeed, both were considering settings where $n \geq d$.

2 PROBLEM SETTINGS AND NOTATIONS

Multivariate linear regression. In this paper, we focus on the multi-output linear regression problem. Suppose that we observe n training examples $D_n = \{(x_1, (y_1^{(1)}, \dots, y_1^{(m)})), \dots, (x_n, (y_n^{(1)}, \dots, y_n^{(m)}))\}$, where $x_i \in \mathbb{R}^d$ and $(y_i^{(1)}, \dots, y_i^{(m)}) \in \mathbb{R}^m$, generated by

$$y_i^{(j)} = \beta_{*j}^\top x_i + \epsilon_i^{(j)}, \quad j = 1, \dots, m,$$

with the unknown signals β_{*j} ($j = 1, \dots, m$) that we want to learn, where $(x_i)_{i=1}^n$ are i.i.d. sample from some distribution P_χ with mean 0 and covariance matrix Σ_x , and $(\epsilon_i^{(j)})_{i=1}^n$ are noise independent of x_i . This setting is called multivariate linear regression, which has been studied for a long time (Breiman and Friedman, 1997; Bai et al., 1992). Indeed, we encounter a multivariate output setting in several situations. Especially, in standard deep learning applications, we typically deal with multivariate output like multi-label classification problems. Our problem setting could be considered as the simplest problem setting to consider such a problem.

Assumption. Throughout the paper, we put the following sub-gaussian assumption on the input random variable. We define $\|Z\|_{\psi_2}$ to be sub-gaussian norm defined by $\|Z\|_{\psi_2} := \inf \{t > 0 : \mathbb{E} [\exp(Z^2/t^2)] \leq 2\}$ for any random variable Z taking its value in \mathbb{R} , and $\|Z\|_{\psi_2} := \sup_{s \neq 0} \|\langle s, Z \rangle / \|s\|\|_{\psi_2}$ for any random variable vector Z taking its value in \mathbb{R}^d . Let design matrix $X := (x_1, \dots, x_n)^\top \in \mathbb{R}^{n \times d}$, and then we assume the following condition on the distribution of X .

Assumption 1. For some positive constants $\sigma_x = O(1)$ and $\sigma = O(1)$, the rows of $X \Sigma_x^{-\frac{1}{2}}$ are sub-gaussian vectors with the sub-gaussian norm at most σ_x and components of $\epsilon^{(i)}$ are independent and have sub-gaussian norms bounded by σ .

Two-layer linear neural network. In this paper, we consider estimating the target linear functions $\beta_{*i}^\top x$ by a two-layer linear neural network which is given by a form of $f_i(x; W, \beta) = \beta_i^\top W x$ where $W \in \mathbb{R}^{d \times d}$ and $\beta_i \in \mathbb{R}^d$ ($i = 1, \dots, m$). When the first-layer's parameter W is given, we estimate the second layer's parameter β_i as the ridge regression estimator:

$$\hat{\beta}_i(W) = W X^\top (X W^\top W X^\top + n \lambda I_n)^{-1} y^{(i)}, \quad i = 1, \dots, m. \quad (2.1)$$

Note that the first layer's parameter W works as a linear coordinate transform. Accordingly, we estimate the target function by $f_i(x; W, \hat{\beta}_i(W)) = \hat{\beta}_i(W)^\top W x$. Then, our question becomes what kind of W yields better performance.

Notations. We write $a \lesssim b$ if there exists an absolute constant c such that $a \leq cb$, and $a \lesssim_{\sigma_x} b$ if there exists an absolute constant c_x that only depends on constant σ_x such that $a \leq c_x b$. For positive semidefinite matrix M , we define $\|\cdot\|_M := \sqrt{x^\top M x}$. For any matrix $M \in \mathbb{R}^{n \times d}$, we define $M_k \in \mathbb{R}^{n \times k}$ to be the matrix which is composed of the first k columns of M , $M_{-k} \in \mathbb{R}^{n \times d-k}$ to be the matrix which is composed of the rest of $d-k$ columns of M . Accordingly, for any matrix $M \in \mathbb{R}^{d \times d}$, we define the eigenvalues of M in decreasing order by $\mu_1(M) \geq \dots \geq \mu_d(M)$, $M_{k,k} \in \mathbb{R}^{k \times k}$ to be the submatrix of M which is composed of the upper left $k \times k$ components

of M , and $M_{-k,-k} \in \mathbb{R}^{d-k \times d-k}$ to be the submatrix of M which is composed of the lower right $d-k \times d-k$ components of M . For matrix A, B , if $B - A$ is positive semidefinite, we write $A \leq B$.

3 MAIN RESULTS

3.1 SELECTING FIRST-LAYER PARAMETER W

Given the ridge estimator for the second layer’s parameter, our interest is what kind of W yields generalization. A natural strategy is choosing W corresponding to vanishingly small predictive risk which is unknown a priori. For this purpose, we aim to construct an estimator of the predictive risk as a function of W . More specifically, we introduce a criterion $R(W)$ as an upper bound of the predictive risk defined as

$$R(W) = \frac{1}{m} \sum_{i=1}^m \min_{\beta_i \in \mathbb{R}^d} \frac{1}{n} \left\| y^{(i)} - XW^\top \beta_i \right\|^2 + \lambda \|\beta_i\|^2 + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right),$$

where $\hat{\Sigma}_{wx} := \frac{1}{n} WX^\top XW^\top$ and σ'^2 is a hyper-parameter. Note that the last term plays a role of regularization on feature mapping W as stated in the introduction. For $\Sigma_{wx} := W\Sigma_x W^\top$, we obtain an upper bound of the predictive risk by bias-variance decomposition as follows:

$$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x^\top \beta_{*i} - x^\top W^\top \hat{\beta}_i(W) \right)^2 \right] \lesssim B + V,$$

where,

$$B := \frac{1}{m} \sum_{i=1}^m \left\| \left(I_d - WX^\top (XW^\top WX^\top + n\lambda I_n)^{-1} XW^\top \right) W^{-\top} \beta_{*i} \right\|_{\Sigma_{wx}}^2,$$

$$V := \frac{1}{m} \sum_{i=1}^m \left\| WX^\top (XW^\top WX^\top + n\lambda I_n)^{-1} \epsilon^{(i)} \right\|_{\Sigma_{wx}}^2.$$

B and V correspond to bias and variance terms respectively. $R(W)$ is essentially important because it indeed yields an upper bound of the bias and variance as stated in the following theorem.

Theorem 1. *There exist (large) constants c_1, c_2 and c_3 which only depend on σ_x such that if t satisfies $t \in (1, \min\{n/c_1, n/4c_1^2\})$, and σ'^2 satisfies $t\sigma^2 \lesssim \sigma'^2$, then it holds that with probability at least $1 - 24e^{-t/c_1} - 8e^{-c_3 n} - \frac{(\log 2)(1+\sigma_x^2)\sigma^2 \max\|\beta_{*i}\|_{\Sigma_x}^2}{n\delta^2}$, for all W which satisfies $R(W) - \sigma^2 \leq \frac{\sigma'^2}{\max\{8c_1^2, c_2\}}$,*

$$B + V \lesssim_{\sigma_x} \max\{R(W) - \sigma^2, \delta\}.$$

Theorem 1 asserts that if we take σ'^2 moderately large, $R(W)$ plays a role of a criterion to select first-layer representation W because $R(W) - \sigma^2$ can be regarded as an estimator for the predictive risk. Therefore, a small function value of $R(W)$ directly leads to small predictive risk, and if we obtain W which makes $R(W) - \sigma^2$ close to zero, the two-layer neural network with such W generalizes in the sense of predictive risk. The regularization term $\sigma'^2 \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right)$ is called “degrees of freedom”, which we can see in Widely Applicable Information Criterion (WAIC) (Watanabe, 2010) or Mallows’ C_p (Mallows, 1973) of ridge regression defined as $\hat{L}(\lambda) := \left\| y - X(X^\top X + n\lambda I_d)^{-1} X^\top y \right\|^2 + 2\hat{\sigma}^2 \text{Tr} \left(\hat{\Sigma}_x \left(\hat{\Sigma}_x + \lambda I_d \right)^{-1} \right)$. This criterion is used to select optimal λ in $d \ll n$ settings. Our proposed function can be regarded as a kind of extension in the sense that we can select optimal basis W in $d \geq n$ settings. Note that, the theoretical properties of Mallows’ C_p and WAIC are valid only in $d \ll n$ settings, and so a theoretical guarantee of the degrees of freedom as a penalty term in high-dimensional settings is a novelty of our work. Typically, in a high dimensional setting where the eigenvalues of Σ_x have slow decay, the ridge regression with positive λ is not necessarily optimal (Wu and Xu) and this kind of classic degrees of freedom may not work because the sum of eigenvalues diverges (Bartlett et al., 2020; Tsigler and Bartlett,

2020). Indeed, diverging eigenvalue-sum is essentially important for obtaining benign overfitting. However, in a two-layer neural network setting, *the coordinate transformation by W can change the problem into like a kernel regime*. The W chosen by $R(W)$ yields an optimal kernel basis in terms of predictive risk, enabling the model to generalize even in situations where it could not generalize without the transformation as we will see later. This can happen if the β_{*j} are well correlated to each other. Such a condition will be distilled to Assumption 3, under which we can show the ridge regression with the optimal W minimizing $R(W)$ gives the Bayes optimal risk.

3.2 EVALUATION OF THE MINIMUM RISK ESTIMATE $R(W) - \sigma^2$

Based on Theorem 1, we may have an upper bound of the predictive risk by evaluating how small $R(W) - \sigma^2$ can be. Here, we give an explicit form of the upper bound by using the eigenvalues of Σ_x and the covariance matrix of $(\beta_{*j})_{j=1}^m$. We let $\Sigma_\beta := \frac{1}{m} \sum_{i=1}^m \beta_{*i} \beta_{*i}^\top$, and define $\sigma_1^2 \geq \dots \geq \sigma_d^2$ and $v_1, \dots, v_d \in \mathbb{R}^d$ as eigenvalues of Σ_β and corresponding eigenvectors respectively. Let $V := (v_1, \dots, v_d)$. Then, we can obtain an upper bound of the infimum of the objective $R(W) - \sigma^2$ as stated in the theorem below.

Theorem 2. *Suppose there exists $k \leq n$ such that $\mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \leq \frac{\sigma'^2}{n} \leq \mu_k \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)$. For $W = \text{diag}(w_1, \dots, w_d) V^\top$ such that $w_i^2 = \frac{n\lambda}{\sigma'^2} \sigma_i^2$, there exists a constant c which only depends on σ_x such that, with probability at least $1 - 4e^{-cn}$,*

$$R(W) - \sigma^2 \lesssim_{\sigma_x} \sum_{i=1}^d \min \left\{ \frac{\sigma'^2}{n}, \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \right\} := U_{\text{NN}}. \quad (3.1)$$

Therefore, as long as the right-hand side of (3.1) is close to 0, W in Theorem 2 can make the predictive risk itself close to 0. This means that the two-layer neural network selected by $R(W)$ can generalize. Qualitatively, this result is due to ‘‘good’’ regularization via feature learning using $R(W)$ on the second-layer parameter. Normal l_2 regularization is often not enough because this regularization can be too strong even in a direction where signals have large contribution, which leads to large bias. In an ideal regularization, directions in which the signals have a large contribution should be weakly regularized, thereby resulting in small bias. Conversely, directions where signals have a small contribution ought to be strongly regularized to facilitate small variance. For this ideal regularization, information about the distribution of signals is essential which is usually not accessible. Thanks to degrees of freedom as a regularization on W , $R(W)$ enables us to make use of the information on the distribution of signals, and this is one of the biggest merits of selecting W by $R(W)$.

3.3 BAYES RISK OPTIMALITY

So far, we have derived an upper bound of the predictive risk. In this section, we show that the derived upper bound can achieve a lower bound under some conditions using the optimal Bayes risk. As stated in the introduction, this is the first attempt to show that an optimized two-layer neural network can achieve Bayes risk optimality. We suppose Σ_β is positive definite. To construct a lower bound of the predictive risk of the two-layer linear neural network we consider here, we introduce the optimal Bayes risk $R_{\text{opt}}(X, \sigma) := \inf_{\hat{\beta}: \text{estimator}} \mathbb{E}_{\beta_* \sim \mathcal{N}(0, \Sigma_\beta)} \left[\mathbb{E}_{Y \sim N(X\beta_*, \sigma^2 I)} \left[\|\beta_* - \hat{\beta}\|_{\Sigma_x}^2 \right] \right]$ where

$X = (x_1, \dots, x_n)$ is a given training input data, $Y = X\beta_* + \epsilon$ with $\epsilon_i \sim N(0, \sigma^2)$, and the infimum is taken over all estimators depending on (X, Y) . The Bayes risk coincides with the predictive risk of the multi-output linear regression averaged over the multiple-output since the Bayes risk is an average of risks over the choice of β_* . Therefore, it is natural to expect that attaining the Bayes optimal risk directly leads to minimizing the predictive risk of the multiple-output regression. The Bayes estimator $\hat{\beta}_{\text{Bayes}}$ is defined as the minimizer that attains the infimum.

Proposition 1. *For positive definite matrices Σ_x, Σ_β , we suppose $x \sim \mathcal{N}(0, \Sigma_x)$, $\beta_* \sim \mathcal{N}(0, \Sigma_\beta)$, $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ and training data $(x_i, y_i)_{i=1}^n \in \mathbb{R}^d \times \mathbb{R}$ is generated by $y_i = \beta^\top x_i + \epsilon_i$. For $X = (x_1, \dots, x_n)$, $\epsilon = (\epsilon_1, \dots, \epsilon_n)^\top$ and $y(\beta_*, \epsilon) = (y_1, \dots, y_n)^\top$, the Bayes estimator $\hat{\beta}_{\text{Bayes}}$ can be explicitly given by*

$$\hat{\beta}_{\text{Bayes}}(\beta_*, \epsilon) = \left(X^\top X + \sigma^2 \Sigma_\beta^{-1} \right)^{-1} X^\top y(\beta_*, \epsilon).$$

This regularization term Σ_β^{-1} enables ideal regularization discussed above. Furthermore, for $W_B = \frac{\sqrt{n\lambda}}{\tilde{\sigma}} \Sigma_\beta^{1/2}$, we can write this optimal Bayes estimator as $\hat{\beta}_{\text{Bayes}} = W_B^\top (W_B X^\top X W_B^\top + nI_d)^{-1} W_B X^\top y$. Therefore, we can represent this Bayes estimator with a two-layer linear neural network and this W_B coincides with W in Theorem 2 except for constant factor. As we will see later, minimum norm interpolator or more generally, normal ridge regression, does not achieve Bayes risk optimality. Asymptotically, this result coincides with the result of (Wu and Xu) when Σ_x and Σ_β have the same eigenvectors. To construct a lower bound of this optimal Bayes risk, for $\tilde{\sigma}^2 = \text{Var} \left[\epsilon_i^{(j)} \right]$ we make the following assumptions:

Assumption 2. We assume that the variance of label noise is not too large, i.e., it holds $2\tilde{\sigma}^2 \leq 1$.

Assumption 3. We assume that the contribution to the output of the components corresponding to small eigenvalues of Σ_x is not too large, that is, there exists $k \leq n$, such that $\sum_{i=k+1}^d \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \leq \tilde{\sigma}^2$.

This assumption ensures that $\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}}$ has a rapidly decreasing eigenvalues, which indicates that alignment between Σ_β and Σ_x has anisotropic property and then we may have better performance by finding the informative direction in the first layer. Under these assumptions, we can construct a lower bound of the predictive risk in the regime of (Tsigler and Bartlett, 2020) as follows:

Theorem 3. Write eigenvectors of $\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}}$ as $u_{B,1}, \dots, u_{B,d} \in \mathbb{R}^d$ and $U_B := (u_{B,1}, \dots, u_{B,d}) \in \mathbb{R}^{d \times d}$. Suppose that components of the rows of $XW_B^\top U_B$ are independent. Then under Assumption 2, 3, for k in Assumption 3 there exists some absolute constant c such that for any t which satisfies $c < t < \min \left\{ n, \frac{n}{2\tilde{\sigma}_x^2} \right\}$ and $k + 2\tilde{\sigma}^2 + \sqrt{kt}\tilde{\sigma}^2 < n/2$, it holds that with probability at least $1 - 24e^{-t/c}$,

$$\min_W \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{x, \epsilon^{(i)}} \left[\left(x^\top \beta_{*i} - x^\top W^\top \hat{\beta}_i(W) \right)^2 \right] \geq R_{\text{opt}}(X, \tilde{\sigma}) \gtrsim \sum_{i=1}^d \min \left\{ \frac{\tilde{\sigma}^2}{n}, \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \right\}.$$

Recalling that the upper bound is $U_{\text{NN}} = \sum_{i=1}^d \min \left\{ \frac{\sigma'^2}{n}, \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \right\}$, Theorem 3 claims that the upper bound of the bias and variance of the optimized two-layer neural network (3.1) achieves the lower bound of the optimal Bayes risk as long as $\sigma'^2 = \Omega(\tilde{\sigma}^2)$. This result characterizes the effectiveness of feature learning using the proposed criterion $R(W)$. In other words, minimizing $R(W)$ can properly find the informative direction, which can be realized thanks to the multiple-output setting. This result indicates that the degrees of freedom is the optimal regularization of feature mapping W in terms of feature learning in the multivariate linear regression problem. Even though $R(W)$ isn't employed in real-world applications, we can establish that in practical terms, neural networks utilize the full potential of feature learning by identifying the implicit regularization that lowers the value of this criterion.

4 COMPARISON WITH NORMAL RIDGE REGRESSION

In this section, we discuss the generalization performance of the two-layer linear neural network whose first-layer representation W is chosen based on $R(W)$, and compare this to a normal ridge estimator without feature learning. Here, we compare the predictive risk of the selected two-layer linear neural network with that of the normal ridge regression (i.e., $W = I$). For that purpose, we derive an upper bound of the predictive risk of the normal ridge regression. This is the first attempt to clarify when a two-layer linear neural network surpasses normal ridge regression in the setting of high-dimensional, multivariate linear regression problems. To construct an upper bound, first, we define eigenvalues of Σ_x and their corresponding eigenvectors as $\lambda_1 \geq \dots \geq \lambda_d$ and $u_1, \dots, u_d \in \mathbb{R}^d$. Accordingly, we define $U := (u_1, \dots, u_d) \in \mathbb{R}^{d \times d}$ and $\tilde{\sigma}_i^2 := u_i^\top \Sigma_\beta u_i$. In the same way as (Tsigler and Bartlett, 2020), for any $k < n$, we define $\rho_k := \frac{1}{n\lambda_{k+1}} (n\lambda + \sum_{i>k} \lambda_i)$ and $A_k(XU, \lambda) := (XU)_{-k}^\top (XU)_{-k} + n\lambda I_n$. Then, we can evaluate the predictive risk of normal ridge regression as stated below.

Corollary 1. Suppose it is known that for some $\delta < 1 - 4^{-n/c_x^2}$, with probability at least $1 - \delta$, the condition number of $A_k(XU, \lambda)$ is at most L . Then there exists a constant c_x , which only depends

on σ_x , such that for any $\lambda \in \mathbb{R}$, $k \leq n/c_x$, and $t \in (1, n/c_x)$, if $A_k(XU, \lambda)$ is positive definite and $\rho_k = O(\frac{1}{L})$, then it holds that with probability at least $1 - \delta - 20e^{-t/c_x}$,

$$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x^\top \beta_{*i} - x^\top \hat{\beta}_i(I_d) \right)^2 \right] \lesssim B_{\text{Norm}} + V_{\text{Norm}},$$

$$\frac{B_{\text{Norm}}}{L^3 \max \left\{ L, \frac{\sum_{i>k} \lambda_i}{n\lambda + \sum_{i>k} \lambda_i} \right\}} = \sum_{i=1}^k \lambda_i \tilde{\sigma}_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \tilde{\sigma}_i^2, \quad (4.1)$$

$$\frac{V_{\text{Norm}}}{L^2} = \frac{kt\sigma^2}{n} + \frac{t\sigma^2}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2}, \quad (4.2)$$

where B_{Norm} corresponds to the bias term and V_{Norm} corresponds to the variance term.

This result is an extension of the result of Tsigler and Bartlett (2020) for multi-output linear regression. k is called effective dimensionality (Bartlett et al., 2020; Tsigler and Bartlett, 2020), which stands for the effective complexity of the ridge regression estimator in high-dimensional settings. Notice that λ in Corollary 1 can be negative. In fact, some negative λ can show higher generalization performance than any $\lambda > 0$ under some conditions (Kobak et al., 2020; Wu and Xu). In particular, when $\lambda = 0$ this ridge estimator coincides with minimum norm interpolator (Bartlett et al., 2020; Liang and Rakhlin, 2020). In a similar way as Tsigler and Bartlett (2020), we can construct lower bounds of the bias B_{Norm} and the variance V_{Norm} (See Lemma 6 and Lemma 7 in Appendix C), and we can also show that the upper bounds and the lower bounds match under some conditions (See Appendix 3). Furthermore, we can demonstrate that the upper bound U_{NN} of the selected model and that of normal ridge regression coincide when Σ_β is isotropic (See Appendix G). This result is natural because when Σ_β is isotropic, the ideal regularization is equivalent to normal l_2 regularization, that is, W selected by the criterion $R(W)$ also becomes close to being isotropic, which indicates that the selected two-layer neural network comes close to normal ridge estimator. In the following part of this section, we discuss the condition that the selected two-layer linear neural network outperforms normal ridge regression comparing U_{NN} and $B_{\text{Norm}} + V_{\text{Norm}}$. Note that since B_{Norm} and V_{Norm} achieve their lower bounds, this comparison is reasonable. For simplicity, we assume that Σ_x and Σ_β have the same eigenvectors.

4.1 WHEN λ_i DECREASE SLOWLY

In this section, we show that there appears separation between feature learning by $R(W)$ and the naive ridge regression when the eigenvalues λ_i decrease slowly. The situation where the decay of λ_i is slow appears in the benign overfitting regime (Bartlett et al., 2020; Xu and Hsu, 2019). The following argument shows selecting W by $R(W)$ is effective in such a setting. Indeed, for the case of normal ridge regression, observing Corollary 1, if λ_i is large in comparison with $\rho_k \lambda_{k+1}$, $\frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2}$ in the first term of the bias B_{Norm} can be seen as a weight that controls the impact of $\lambda_i \sigma_i^2$ to the bias. However, if the decay of λ_i is gradual, the weight $\frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2}$ can be close to 1, thereby enhancing the effect of the bias. On the other hand, in the context of the two-layer linear neural network which is chosen based on $R(W)$, the upper bound of Theorem 2 is written as $U_{\text{NN}} = \sum_{i=1}^d \min \left\{ \frac{\sigma_i'^2}{n}, \lambda_i \sigma_i^2 \right\}$.

This bound implies that the effect of bias $\lambda_i \sigma_i^2$ is no larger than $\frac{\sigma_i'^2}{n}$ and this prevents the bias term from deteriorating the predictive risk of the model. The example below reflects the difference in the generalization performance of both models discussed above.

Example 1. We suppose $\sqrt{n} \in \mathbb{N}$ and eigenvalues of Σ_x decay slowly, that is, we consider the situation as follows:

$$\lambda_i = \begin{cases} 1 & i \leq n \\ \frac{1}{i+1-n} & i \geq n \end{cases}, \quad \sigma_i = \frac{n}{\log n} e^{-i}, \quad \sigma^2 = 1, \quad \sigma_i'^2 = 1, \quad d \geq 2n.$$

Then for any $k < n$ in Corollary 1, it holds that

$$B_{\text{Norm}} \gtrsim 1, \quad U_{\text{NN}} \lesssim \frac{\log n}{n}.$$

Example 1 claims that the two-layer linear neural network selected by $R(W)$ can generalize even when the decay of Σ_x is slow and normal ridge regression can't generalize. This result reflects the

effect of the good regularization discussed above, using W obtained by feature learning via $R(W)$. This allows for weak regularization in directions where signals make significant contributions, and strong regularization where the signals' contributions are minor.

4.2 WHEN x AND β_{*i} ARE MISALIGNED

In this section, we compare the predictive performance of the selected two-layer neural network and normal ridge regression when x and β_{*i} are misaligned. First, we introduce the notion of alignment and misalignment. We say that x and β_{*i} are aligned when the directions of large variance for x and that of β_{*i} closely match. In other words, an aligned situation is that the eigenvectors of Σ_x corresponding to large eigenvalues and those of Σ_β are in close directions. On the other hand, if the direction in which x has large variance and that of β_{*i} are different, that is, if eigenvectors of Σ_x corresponding to large eigenvalues and those of Σ_β are orthogonal, we refer to x and β_{*i} as being misaligned. We may define the alignment in a more rigorous way by considering the situation where Σ_x and Σ_β have the same eigenvectors. In this setting, for $\lambda_1 \geq \dots \geq \lambda_d$ if $\sigma_1 \geq \dots \geq \sigma_d$ then we say x and β_{*i} are aligned, and on the contrary, if $\sigma_1^2 \leq \dots \leq \sigma_d^2$, then we say x and β_{*i} are misaligned. Here, we consider a misaligned situation, that is, $\sigma_1^2 \leq \dots \leq \sigma_d^2$. For the case of normal ridge regression, for any $k < n$, we can write the bias term as $B_{\text{Norm}} = \sum_{i=1}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2$. Since $\sigma_1^2 \leq \dots \leq \sigma_d^2$, $\sigma_{k+1}^2, \dots, \sigma_d^2$ can be large and this leads to large bias. On the other hand, the two-layer neural network selected by $R(W)$ can decrease both bias and variance thanks to good regularization via learning with $R(W)$ as the example below demonstrates.

Example 2. We suppose x and β_{*i} are misaligned, that is, we consider the situation as follows:

$$\lambda_i = i^{-1}, \sigma_i^2 = i e^{i-d}, \sigma^2 = 1, \sigma'^2 = 1.$$

Then for any $k < n$ in Corollary 1, it holds that

$$B_{\text{Norm}} \gtrsim 1, \quad U_{\text{NN}} \lesssim \frac{\log n}{n}.$$

This result indicates that the estimator selected by $R(W)$ outperforms the normal ridge estimator in misaligned situations.

4.3 WHEN TAIL EIGENVALUES OF Σ_x ARE LARGE

In this section, we show that the selected two-layer linear neural network outperforms normal ridge regression when tail eigenvalues of Σ_x are large. In normal ridge regression, for any $k = o(n)$, the variance term is written as $\frac{V_{\text{Norm}}}{L^2 t \sigma^2} = \frac{k}{n} + \frac{1}{n \rho_k^2 \lambda_{k+1}^2} \sum_{i=k+1}^d \lambda_i^2$. We can see that if tail eigenvalues of Σ_x are large, that is, if $\sum_{i>k} \lambda_i^2$ is large, then the second term of V_{Norm} can also be large. This indicates that normal ridge regression doesn't generalize. On the other hand, because of feature learning via $R(W)$, the selected two-layer linear neural network can generalize even when the normal ridge regression can't generalize as the following example shows.

Example 3. Suppose that tail eigenvalues of Σ_x are large, that is, we consider the situation as follows:

$$\lambda_i = \begin{cases} i^{-1} & i \leq n \\ \frac{1}{n} & i > n \end{cases}, \sigma_i^2 = i^{-2}, \sigma^2 = 1, \sigma'^2 = 1, d = n^3.$$

Then for any $k < n$ in Corollary 1, it holds that

$$B_{\text{Norm}} + V_{\text{Norm}} \gtrsim 1, \quad U_{\text{NN}} \lesssim n^{-\frac{2}{3}}.$$

Qualitatively, this result also reflects the good regularization obtained by $R(W)$. W selected by $R(W)$ enables strong regularization in directions where signals don't contribute significantly, leading to a decrease in variance without a drastic increase in bias.

4.4 WHEN $y^{(i)}$ IS LARGE

In this section, we demonstrate that the selected two-layer linear neural network surpasses normal ridge regression when the norm of $y^{(i)}$ is notably large. In normal ridge regression, the bias can be written as $B_{\text{Norm}} = \sum_{i=1}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2$ and the weights $\frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2}$ in the first term contribute to reducing the impact on the bias term. However, if $y^{(i)}$ is large, that is, $\lambda_i \sigma_i^2$ is large, then

the bias term doesn't vanish even if the weight is small. This means that an isotropic regularization is not enough to effectively decrease the effect of bias in such situations. On the contrary, the selected model with good regularization via feature learning by $R(W)$ is capable of bringing both the bias and variance close to zero. The following example demonstrates the difference between the selected two-layer neural network and normal ridge regression.

Example 4. We suppose $n^{\frac{2}{3}} \in \mathbb{N}$ and norm of $y^{(i)}$ is large, that is, we consider the situation as follows:

$$\lambda_i = i^{-1}, \sigma_i^2 = \begin{cases} n^{\frac{2}{3}} & i \leq n^{\frac{2}{3}} \\ i^{-1} & i > n^{\frac{2}{3}} \end{cases}, \sigma^2 = 1, \sigma'^2 = 1.$$

Then for any $k < n$ in Corollary 1, it holds that

$$B_{\text{Norm}} + V_{\text{Norm}} \gtrsim 1, \quad U_{\text{NN}} \lesssim n^{-\frac{1}{3}}.$$

5 NUMERICAL EXPERIMENTS

In this section, we conduct numerical experiments to justify our theoretical results. We consider the setting that $(d, n, m) = (1, 000, 500, 1, 000)$, $\Sigma_x = \text{diag}(1, 2^{-1}, \dots, 1, 000^{-1})$, $x \sim \mathcal{N}(0, \Sigma_x)$, $\Sigma_\beta = \text{diag}(1, 2^{-a}, \dots, 1, 000^{-a})$ ($a = 0, 0.5, 1, \dots, 10$), $(\sigma^2, \sigma'^2) = (1, 2)$ and $\epsilon_i^{(j)} \sim_{\text{i.i.d.}} \mathcal{N}(0, \sigma^2)$. Since the criterion $R(W)$ is non-convex, we optimized this function with gradient Langevin dynamics, which is a noisy gradient descent method aiming to escape from sharp local minima. For a random variable matrix ξ_t which satisfies $(\xi_t)_{i,j} \sim_{\text{i.i.d.}} \mathcal{N}(0, 1)$, learning rate η and inverse temperature parameter β , the update of the gradient Langevin dynamics can be written as $W^{(t+1)} = W^{(t)} - \eta \nabla R(W^{(t)}) + \sqrt{\frac{2\eta}{\beta}} \xi_t$. In this experiment, we set $(\eta, \beta) = (0.1, 100, 000)$ and iterated the update 500 times and selected W corresponding to the smallest $R(W)$. Figure 1 shows the comparison of the predictive risks with a base-10 logarithm between the obtained two-layer linear neural network, normal ridge regression, and Bayes-optimal estimator. The regularization parameters for ridge regression were chosen from 1.0, 0.5, 0.1, 0.01, 0.001, 0.0001 to minimize the predictive risk calculated using additional data. When a is small and the eigenvalues of Σ_β decay slowly, the proposed method has equivalent performance to or falls below the ridge regression. However, as a becomes larger and the eigenvalues of Σ_β decay more quickly, it can be seen that the proposed method begins to outperform ridge regression. This reflects the fact that the more isotropic Σ_β is, the more isotropic the optimal parameter of the first layer becomes, hence the benefit of finding informative directions by feature learning is diminished in such cases. 2 shows the minimum values of the proposed criterion $R(W) - \sigma^2$ with a base-10 logarithm obtained by gradient Langevin dynamics and predictive risks with a base-10 logarithm of the two-layer linear neural network with corresponding W . We can see that the behavior of the function values of $R(W) - \sigma^2$ captures that of predictive risk. This result shows that the proposed criterion $R(W) - \sigma^2$ can play a role of an estimator for an upper bound of the predictive risk.




Figure 1: Comparison of predictive risks.




Figure 2: $R(W) - \sigma^2$ and predictive risk.

6 CONCLUSION AND FUTURE WORK

In this paper, we have considered the feature learning of a multi-output linear regression problem using a two-layer linear neural network. We have proposed a criterion with penalty term on feature mapping W like Mallows' C_p and have shown that an estimator selected by the criterion leads to making its predictive risk close to zero. Then, we have shown the optimized model by the proposed criterion can achieve nearly optimal Bayes risk. Furthermore, we have characterized when normal

ridge regression doesn't generalize and shown that the selected two-layer linear neural network generalizes even in such situations. We state the future directions of this research. On the optimization aspect, we have to guarantee that we can minimize the proposed criterion well. It is not obvious whether we can obtain the optimal solution by standard optimization techniques such as stochastic gradient descent because the objective is non-convex. Another future direction is replacing the activation function. Although we considered identity in this paper, we usually use nonlinear activation functions such as ReLU in real applications.

7 ACKNOWLEDGEMENT

KS was partially supported by JST CREST (JPMJCR2015). TS was partially supported by JSPS KAKENHI (20H00576) and JST CREST (JPMJCR2115).

REFERENCES

- S. Azulay, E. Moroshko, M. S. Nacson, B. E. Woodworth, N. Srebro, A. Globerson, and D. Soudry. On the implicit bias of initialization shape: Beyond infinitesimal mirror descent. In *International Conference on Machine Learning*, pages 468–477. PMLR, 2021.
- J. Ba, M. A. Erdogdu, T. Suzuki, Z. Wang, D. Wu, and G. Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In *Advances in Neural Information Processing Systems*, 2022.
- Z. Bai, C. R. Rao, and Y. Wu. M-estimation of multivariate linear regression parameters under a convex discrepancy function. *Statistica Sinica*, 2(1):237–254, 1992.
- P. L. Bartlett, P. M. Long, G. Lugosi, and A. Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- M. Belkin, S. Ma, and S. Mandal. To understand deep learning we need to understand kernel learning. In *International Conference on Machine Learning*, pages 541–549. PMLR, 2018.
- M. Belkin, A. Rakhlin, and A. B. Tsybakov. Does data interpolation contradict statistical optimality? In *International Conference on Artificial Intelligence and Statistics*, pages 1611–1619. PMLR, 2019.
- L. Breiman and J. H. Friedman. Predicting multivariate responses in multiple linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(1):3–54, 1997.
- Y. Cao, Z. Chen, M. Belkin, and Q. Gu. Benign overfitting in two-layer convolutional neural networks. *arXiv preprint arXiv:2202.06526*, 2022.
- N. S. Chatterji and P. M. Long. Finite-sample analysis of interpolating linear classifiers in the overparameterized regime. *Journal of Machine Learning Research*, 22:129–1, 2021.
- N. S. Chatterji, P. M. Long, and P. L. Bartlett. The interplay between implicit bias and benign overfitting in two-layer linear networks. *Journal of Machine Learning Research*, 23(263):1–48, 2022.
- A. Damian, J. Lee, and M. Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pages 5413–5452. PMLR, 2022.
- E. Dobriban and S. Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.
- S. Frei, N. S. Chatterji, and P. Bartlett. Benign overfitting without linearity: Neural network classifiers trained by gradient descent for noisy linear data. In *Conference on Learning Theory*, pages 2668–2703. PMLR, 2022.
- T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986, 2022.

- D. Kobak, J. Lomond, and B. Sanchez. The optimal ridge penalty for real-world high-dimensional data can be zero or negative due to the implicit ridge regularization. *The Journal of Machine Learning Research*, 21(1):6863–6878, 2020.
- Z. Li, W. J. Su, and D. Sejdinovic. Benign overfitting and noisy features. *Journal of the American Statistical Association*, pages 1–13, 2022.
- T. Liang and A. Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize. *The Annals of Statistics*, 48(3):1329–1347, 2020.
- C. L. Mallows. Some Comments on C_p . *Technometrics*, 15(4):661–675, 1973.
- S. Mei and A. Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766, 2022.
- A. Montanari, F. Ruan, Y. Sohn, and J. Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*, 2019.
- V. Muthukumar, K. Vodrahalli, V. Subramanian, and A. Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 1(1):67–83, 2020.
- J. Negrea, G. K. Dziugaite, and D. Roy. In defense of uniform convergence: Generalization via derandomization with an application to interpolating predictors. In *International Conference on Machine Learning*, pages 7263–7272. PMLR, 2020.
- A. Tsigler and P. L. Bartlett. Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286*, 2020.
- R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing: Theory and Applications*, pages 210–268. Cambridge University Press, 2012.
- S. Watanabe. Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11:3571–3594, 2010.
- D. Wu and J. Xu. On the optimal weighted l_2 regularization in overparameterized linear regression. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, pages 10112–10123. Curran Associates, Inc.
- J. Xu and D. J. Hsu. On the number of variables to use in principal component regression. *Advances in Neural Information Processing Systems*, 32, 2019.

A PROOF OF THEOREM 1

Proof. We define $X_w = XW^\top$, $\Sigma_{wx} = W\Sigma_xW^\top$, $\hat{\Sigma}_{wx} = W\hat{\Sigma}_xW^\top$ and $\Sigma_{w\beta} = W^{-\top}\Sigma_\beta W^{-1}$. Accordingly, we introduce the following notations:

- Let $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_d$ be eigenvalues of $\hat{\Sigma}_{wx}$ and $\hat{u}_1, \dots, \hat{u}_d$ be corresponding eigenvectors. Then, we define $\hat{U} = (\hat{u}_1, \dots, \hat{u}_d)$.
- Let $\hat{\Sigma}_{\hat{u}x} := \hat{U}^\top \hat{\Sigma}_{wx} \hat{U} = \text{diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_d)$.
- Let $\Sigma_{\hat{u}x} := \hat{U}^\top \Sigma_{wx} \hat{U}$.
- Let $\beta_{wi} = W^{-\top} \beta_{*i}$ and $\Sigma_{w\beta} = W^{-\top} \Sigma_\beta W^{-1}$.
- Let $\hat{\sigma}_i^2 := \hat{u}_i^\top \Sigma_{w\beta} \hat{u}_i$, and $\Sigma_{\hat{u}\beta} := \hat{U}^\top \Sigma_{w\beta} \hat{U}$.
- Let $A_k(X_w \hat{U}, \lambda) = \left(X_w \hat{U} \right)_{-k} \left(X_w \hat{U} \right)_{-k}^\top + n\lambda I_n$

From Proposition 2, there exists a constant c_x which only depends on σ_x such that for any t and k which satisfies $t \in (1, n/c_x)$ and $\sqrt{k} + \sqrt{t} \leq \sqrt{n}/c_x$, if $\text{cond}\left(A_k(X_w \hat{U}, \lambda)\right) \leq L$, then it holds that with probability at least $1 - 20e^{-t/c_x}$,

$$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x^\top \beta_{*i} - x^\top W^\top \hat{\beta}_i(W, \lambda) \right)^2 \right] \lesssim_{\sigma_x} B + V,$$

$$\begin{aligned} \frac{B}{L^4} &= \text{Tr} \left((\Sigma_{\hat{u}x})_{k,k}^{-1} (\Sigma_{\hat{u}\beta})_{-k,-k} \right) \left(\lambda + \frac{1}{n} \text{Tr} \left((\Sigma_{\hat{u}x})_{-k,-k} \right) \right)^2 \\ &\quad + \text{Tr} \left((\Sigma_{\hat{u}x})_{-k,-k} (\Sigma_{\hat{u}\beta})_{-k,-k} \right), \end{aligned} \quad (\text{A.1})$$

$$\frac{V}{L^2 t \sigma^2} = \frac{k}{n} + \frac{1}{n} \frac{\text{Tr} \left((\Sigma_{\hat{u}x})_{-k,-k}^2 \right)}{\left(\lambda + \frac{1}{n} \text{Tr} \left((\Sigma_{\hat{u}x})_{-k,-k} \right) \right)^2}. \quad (\text{A.2})$$

From Lemma 2, there exist constant c_b, C_b which only depend on σ_x such that it holds with probability at least $1 - 2e^{-c_b n}$ for any $k' < n$,

$$\begin{aligned} &\left(\frac{3}{2} + C_b \sqrt{\frac{k'}{n}} \right)^{-1} \left(\text{Tr} \left(\left((\Sigma_{\hat{u}x})_{k',k'} + \lambda I'_k \right)^{-1} (\Sigma_{\hat{u}x})_{k',k'} (\Sigma_{\hat{u}\beta})_{k',k'} \right) \right) \\ &\leq \text{Tr} \left(\left(\left(\hat{\Sigma}_{\hat{u}x} \right)_{k',k'} + \lambda I'_k \right)^{-1} \left(\hat{\Sigma}_{\hat{u}x} \right)_{k',k'} (\Sigma_{\hat{u}\beta})_{k',k'} \right) \\ &\leq \left(\frac{3}{2} + C_b \sqrt{\frac{k'}{n}} \right) \left(\text{Tr} \left(\left((\Sigma_{\hat{u}x})_{k',k'} + \lambda I'_k \right)^{-1} (\Sigma_{\hat{u}x})_{k',k'} (\Sigma_{\hat{u}\beta})_{k',k'} \right) \right). \end{aligned} \quad (\text{A.3})$$

For these c_x and C_b , we assume that it holds $R(W) - \sigma^2 \leq \frac{\sigma'^2}{\max\{8c_x^2, 8C_b^2\}}$. Then, we define k_w such that $\mu_{k_w}(\hat{\Sigma}_{wx}) \leq \lambda \leq \mu_{k_w+1}(\hat{\Sigma}_{wx})$. Suppose $k_w > \frac{n}{\max\{4c_x^2, 4C_b^2\}}$, then

$$\begin{aligned} R(W) - \sigma^2 &\geq \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\ &= \frac{\sigma'^2}{n} \sum_{i=1}^n \frac{\mu_i(\hat{\Sigma}_{wx})}{\lambda + \mu_i(\hat{\Sigma}_{wx})} \\ &\geq \frac{\sigma'^2}{n} \sum_{i=1}^{k_w} \frac{\mu_i(\hat{\Sigma}_{wx})}{\lambda + \mu_i(\hat{\Sigma}_{wx})} \\ &\geq \frac{\sigma'^2}{2n} k_w \\ &> \frac{\sigma'^2}{\max\{8c_x^2, 8C_b^2\}}. \end{aligned}$$

This contradicts with the assumption above, and therefore it holds $k_w \leq \frac{n}{\max\{4c_x^2, 4C_b^2\}}$. We set $k = k_w$ and $k' = k_w$. Since $k_w \leq \frac{n}{4c_x^2}$, if t satisfies $\sqrt{n}/2c_x + \sqrt{t} \leq \sqrt{n}/c_x$, then t also satisfies $\sqrt{k_w} + \sqrt{t} \leq \sqrt{n}/c_x$. Combining with $t \in (1, n/c_x)$, the condition on t can be written as $t \in (1, \min\{n/c_x, n/4c_x^2\})$. Furthermore, since $k_w \leq \frac{n}{4C_b^2}$, we can also evaluate as

$$\begin{aligned} &\frac{1}{2} \left(\text{Tr} \left(\left((\Sigma_{\hat{u}x})_{k_w, k_w} + \lambda I'_k \right)^{-1} (\Sigma_{\hat{u}x})_{k_w, k_w} (\Sigma_{\hat{u}\beta})_{k_w, k_w} \right) \right) \\ &\leq \text{Tr} \left(\left(\left(\hat{\Sigma}_{\hat{u}x} \right)_{k_w, k_w} + \lambda I'_k \right)^{-1} \left(\hat{\Sigma}_{\hat{u}x} \right)_{k_w, k_w} (\Sigma_{\hat{u}\beta})_{k_w, k_w} \right) \end{aligned}$$

$$\leq 2 \left(\text{Tr} \left(\left((\Sigma \hat{u}x)_{k_w, k_w} + \lambda I_k \right)^{-1} (\Sigma \hat{u}x)_{k_w, k_w} (\Sigma \hat{u}\beta)_{k_w, k_w} \right) \right).$$

From Lemma 3, there exists constant c_1 such that with probability at least $1 - 4e^{-c_1 n}$, it holds

$$\lambda + \frac{1}{n} \text{Tr} \left((\Sigma \hat{u}x)_{-k_w, -k_w} \right) \lesssim \lambda \lesssim \mu_k \left((\Sigma \hat{u}x)_{k_w, k_w} \right).$$

Therefore, we can evaluate the first term of A.1 as

$$\begin{aligned} & \text{Tr} \left((\Sigma \hat{u}x)_{k_w, k_w}^{-1} (\Sigma \hat{u}\beta)_{k_w, k_w} \right) \left(\lambda + \frac{1}{n} \text{Tr} \left((\Sigma \hat{u}x)_{-k_w, -k_w} \right) \right)^2 \\ & \lesssim \lambda^2 \text{Tr} \left((\Sigma \hat{u}x)_{k_w, k_w}^{-1} (\Sigma \hat{u}\beta)_{k_w, k_w} \right) \\ & = \sum_{i=1}^{k_w} \mu_i \left((\Sigma \hat{u}x)_{k_w, k_w} \right) \hat{\sigma}_{k_w, i}^2 \frac{\lambda^2}{\mu_i \left((\Sigma \hat{u}x)_{k_w, k_w} \right)^2} \\ & \lesssim \sum_{i=1}^{k_w} \mu_i \left((\Sigma \hat{u}x)_{k_w, k_w} \right) \hat{\sigma}_{k_w, i}^2 \frac{\lambda^2}{\left(\mu_i \left((\Sigma \hat{u}x)_{k_w, k_w} \right) + \lambda \right)^2} \\ & \leq \sum_{i=1}^{k_w} \mu_i \left((\Sigma \hat{u}x)_{k_w, k_w} \right) \hat{\sigma}_{k_w, i}^2 \frac{\lambda}{\mu_i \left((\Sigma \hat{u}x)_{k_w, k_w} \right) + \lambda} \\ & = \lambda \text{Tr} \left(\left((\Sigma \hat{u}x)_{k_w, k_w} + \lambda I_{k_w} \right)^{-1} (\Sigma \hat{u}x)_{k_w, k_w} (\Sigma \hat{u}\beta)_{k_w, k_w} \right). \end{aligned}$$

Therefore, we can evaluate (A.1) as

$$\begin{aligned} \frac{B}{L^4} & \lesssim \lambda \text{Tr} \left(\left((\Sigma \hat{u}x)_{k_w, k_w} + \lambda I_{k_w} \right)^{-1} (\Sigma \hat{u}x)_{k_w, k_w} (\Sigma \hat{u}\beta)_{k_w, k_w} \right) \\ & \quad + \text{Tr} \left((\Sigma \hat{u}x)_{-k_w, -k_w} (\Sigma \hat{u}\beta)_{-k_w, -k_w} \right). \end{aligned}$$

By Lemma 11, it holds that with probability at least $1 - 2e^{-c_2 t}$,

$$\begin{aligned} & \text{Tr} \left(\left(\hat{\Sigma} \hat{u}x \right)_{-k_w, -k_w} (\Sigma \hat{u}\beta)_{-k_w, -k_w} \right) \\ & \geq \left(1 - \frac{t}{n} \sigma_x^2 \right) \text{Tr} \left((\Sigma \hat{u}x)_{-k_w, -k_w} (\Sigma \hat{u}\beta)_{-k_w, -k_w} \right). \end{aligned}$$

Combining (A) and these two inequalities, we have

$$\begin{aligned} \frac{B}{L^4} & \lesssim_{\sigma_x} \text{Tr} \left(\left(\left(\hat{\Sigma} \hat{u}x \right)_{k_w, k_w} + \lambda I_{k_w} \right)^{-1} \left(\hat{\Sigma} \hat{u}x \right)_{k_w, k_w} (\Sigma \hat{u}\beta)_{k_w, k_w} \right) \\ & \quad + \text{Tr} \left(\left(\hat{\Sigma} \hat{u}x \right)_{-k_w, -k_w} (\Sigma \hat{u}\beta)_{-k_w, -k_w} \right). \end{aligned}$$

By the way, since $\hat{\lambda}_{k_w+1} \leq \lambda \leq \hat{\lambda}_{k_w}$ by the definition of k_w , we can evaluate as follows:

$$\begin{aligned} & \frac{\lambda}{m} \sum_{i=1}^m \beta_{wi}^\top X_w^\top (X_w X_w^\top + n \lambda I_n)^{-1} X_w \beta_{wi} \\ & = \lambda \text{Tr} \left(\left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \hat{\Sigma}_{wx} \Sigma_w \beta \right) \\ & = \sum_{i=1}^d \frac{\lambda}{\lambda + \hat{\lambda}_i} \hat{\lambda}_i \hat{\sigma}_i^2 \\ & \gtrsim \sum_{i=1}^{k_w} \frac{\lambda}{\lambda + \hat{\lambda}_i} \hat{\lambda}_i \hat{\sigma}_i^2 + \sum_{i=k_w+1}^d \hat{\lambda}_i \hat{\sigma}_i^2 \end{aligned}$$

$$\begin{aligned}
&= \lambda \text{Tr} \left(\left(\left(\hat{\Sigma}_{\hat{u}x} \right)_{k_w, k_w} + \lambda I_{k_w} \right)^{-1} \left(\hat{\Sigma}_{\hat{u}x} \right)_{k_w, k_w} (\Sigma_{\hat{u}\beta})_{k_w, k_w} \right) \\
&\quad + \text{Tr} \left(\left(\hat{\Sigma}_{\hat{u}x} \right)_{-k_w, -k_w} (\Sigma_{\hat{u}\beta})_{-k_w, -k_w} \right).
\end{aligned}$$

Combining these two inequalities, we have

$$\frac{B}{L^4} \lesssim_{\sigma_x} \frac{\lambda}{m} \sum_{i=1}^m \beta_{wi}^\top X_w^\top (X_w X_w^\top + n\lambda I_n)^{-1} X_w \beta_{wi}. \quad (\text{A.4})$$

For the variance term (A.2), since $\lambda > 0$ and $\hat{\lambda}_{k_w+1} \leq \lambda$, we can evaluate as

$$\begin{aligned}
\frac{V}{L^2 t \sigma^2} &= \frac{k_w}{n} + \frac{1}{n} \frac{\text{Tr} \left((\Sigma_{\hat{u}x})_{-k_w, -k_w} \left(\hat{\Sigma}_{\hat{u}x} \right)_{-k_w, -k_w} \right)}{\left(\lambda + \frac{1}{n} \text{Tr} \left((\Sigma_{\hat{u}x})_{-k_w, -k_w} \right) \right)^2} \\
&\leq \frac{k_w}{n} + \frac{\lambda}{n\lambda^2} \text{Tr} \left((\Sigma_{\hat{u}x})_{-k_w, -k_w} \right) \\
&= \frac{k_w}{n} + \frac{1}{n\lambda} \text{Tr} \left((\Sigma_{\hat{u}x})_{-k_w, -k_w} \right).
\end{aligned}$$

In the same way as the bias, we can evaluate the variance as follows:

$$\begin{aligned}
\frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) &= \frac{\sigma'^2}{n} \sum_{i=1}^d \frac{\hat{\lambda}_i}{\hat{\lambda}_i + \lambda} \\
&\gtrsim \frac{k_w \sigma'^2}{n} + \frac{\sigma'^2}{n} \sum_{i=k_w+1}^d \frac{\hat{\lambda}_i}{\hat{\lambda}_i + \lambda} \\
&\geq \frac{k_w \sigma'^2}{n} + \frac{\sigma'^2}{n\lambda} \sum_{i=k_w+1}^d \hat{\lambda}_i \\
&= \frac{k_w \sigma'^2}{n} + \frac{\sigma'^2}{n\lambda} \text{Tr} \left(\left(\hat{\Sigma}_{\hat{u}x} \right)_{-k_w, -k_w} \right).
\end{aligned}$$

By Lemma 11, it holds that with probability at least $1 - 2e^{-c_3 t}$,

$$\text{Tr} \left(\left(\hat{\Sigma}_{\hat{u}x} \right)_{-k_w, -k_w} \right) \geq \left(1 - \frac{t}{n} \sigma_x^2 \right) \text{Tr} \left((\Sigma_{\hat{u}x})_{-k_w, -k_w} \right).$$

Combining these three inequalities, if $t\sigma^2 \lesssim \sigma'^2$, we have

$$\frac{V}{L^2 t \sigma^2} \lesssim_{\sigma_x} \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right). \quad (\text{A.5})$$

Let $\bar{U} \in \mathbb{R}^{n \times n}$ be eigenvectors of $X_w X_w^\top$. From Lemma 1 and Lemma 13, for any $t \in (0, n)$ it holds that with probability at least $1 - \frac{(\log 2)(1 + \sigma_x^2)\sigma^2 \max\|\beta_{*i}\|_{\Sigma_x}^2}{n\delta^2} - 2e^{-c_4 t} - 2e^{-c_5 n}$,

$$\begin{aligned}
&R(W) \\
&= \frac{1}{m} \sum_{i=1}^m \min_{\beta \in \mathbb{R}^t} \frac{1}{n} \left\| \mathbf{y}^{(i)} - X_w \beta \right\|^2 + \lambda \|\beta\|^2 + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
&= \frac{\lambda}{m} \sum_{i=1}^m \mathbf{y}^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} \mathbf{y}^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
&= \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i}
\end{aligned}$$

$$\begin{aligned}
& + \frac{\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
= & \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \\
& + \frac{\lambda}{nm} \sum_{i=1}^m \epsilon^{(i)\top} \left(\hat{\Sigma}_{wx} + \lambda I_n \right)^{-1} \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
= & \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \\
& + \frac{1}{nm} \sum_{i=1}^m \sum_{j=1}^d \frac{\lambda}{\lambda + \hat{\lambda}_j} \epsilon^{(i)\top} \bar{U}_j^\top \bar{U}_j \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
\geq & \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \\
& + \frac{1}{nm} \sum_{i=1}^m \sum_{j=k_w+1}^d \frac{\lambda}{\lambda + \hat{\lambda}_j} \epsilon^{(i)\top} \bar{U}_j^\top \bar{U}_j \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
\gtrsim & \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \\
& + \frac{1}{nm} \sum_{i=1}^m \sum_{j=k_w+1}^d \epsilon^{(i)\top} \bar{U}_j^\top \bar{U}_j \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
\gtrsim & \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \\
& + \sigma^2 + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
\geq & \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\
& + \sigma^2 - 2\delta.
\end{aligned}$$

Therefore, if $t\sigma^2 \lesssim \sigma'^2$, combining with A.4 and A.5 we have

$$\frac{B}{L^4} + \frac{V}{L^2} \lesssim_{\sigma_x} \max\{R(W) - \sigma^2, \delta\}.$$

Finally, we need to check $\text{cond} \left(A_{k_w} \left(XW^\top \hat{U}, \lambda \right) \right) = \Theta(1)$. Since

$$\frac{1}{n} \mu_{\max} \left(X_w \hat{U}_{-k_w}^\top \hat{U}_{-k_w} X_w^\top \right) = \mu_{\max} \left(\hat{\Sigma}_{\hat{u}x} \right) = \frac{1}{n} S_{\max} \left(X_w \hat{U}_{-k_w}^\top \right)^2 \leq \lambda,$$

we have

$$\text{cond} \left(A_{k_w} \left(XW^\top \hat{U}, \lambda \right) \right) \leq \frac{n\lambda + n\mu_{\max} \left(X_w \hat{U}_{-k_w}^\top \hat{U}_{-k_w} X_w^\top \right)}{n\lambda} \leq \frac{n\lambda + n\lambda}{n\lambda} \leq 2 (= L).$$

Therefore, it holds

$$B + V \lesssim_{\sigma_x} \max\{R(W) - \sigma^2, \delta\}.$$

□

In the later part of this section, we show some lemmas needed to prove Theorem 1.

Lemma 1. *In the setting of Theorem 1, there exists a constant $c > 0$ which only depends on σ_x such that with probability at least $1 - \frac{(\log 2)(1+\sigma_x^2)\sigma^2 \max_i \|\beta_{*i}\|_{\Sigma_x}^2}{n\delta^2} - 2e^{-cn}$,*

$$\frac{\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} < \delta.$$

Proof. For any i , $\epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i}$ is mean 0 and variance $\text{Var} \left[\epsilon_1^{(1)} \right] \left\| (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \right\|^2$ random variable. We define index $I(W, \lambda)$ as follows:

$$I(W, \lambda) = \underset{i}{\text{argmin}} \left| \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \right|.$$

Since $\epsilon_j^{(i)}$ are sub-gaussian with the sub-gaussian norm at most σ , $\text{Var} \left[\epsilon_j^{(i)} \right] \leq (\log 2)\sigma^2$. By Chebyshev's inequality, we obtain

$$\begin{aligned} & \Pr \left(\left| \left(\epsilon^{I(W, \lambda)} \right)^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*I(W, \lambda)} \right| < \delta \right) \\ & > 1 - \frac{\text{Var} \left[\epsilon_1^{(1)} \right] \left\| (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*I(W, \lambda)} \right\|^2}{\delta^2} \\ & > 1 - \frac{\text{Var} \left[\epsilon_1^{(1)} \right] \left\| X \beta_{*I(W, \lambda)} \right\|^2}{n^2 \lambda^2 \delta^2} \\ & > 1 - \frac{(\log 2)\sigma^2 \left\| X \beta_{*I(W, \lambda)} \right\|^2}{n^2 \lambda^2 \delta^2}. \end{aligned}$$

From Lemma 11, there exists a constant c such that it holds with probability at least $1 - 2e^{-cn}$,

$$\begin{aligned} & \Pr \left(\frac{1}{n} \left\| X \beta_{*I(W, \lambda)} \right\|^2 < (1 + \sigma_x^2) \max_i \|\beta_{*i}\|_{\Sigma_x}^2 \right) \\ & > \Pr \left(\frac{1}{n} \left\| X \beta_{*I(W, \lambda)} \right\|^2 < \left(1 + \frac{t}{n} \sigma_x^2 \right) \left\| \beta_{*I(W, \lambda)} \right\|_{\Sigma_x}^2 \right) \\ & \geq 1 - 2e^{-ct}. \end{aligned}$$

Combining these inequalities, it holds that with probability at least $1 - 2e^{-cn}$,

$$\begin{aligned} & \Pr \left(\left| \left(\epsilon^{I(W, \lambda)} \right)^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*I(W, \lambda)} \right| < \delta \right) \\ & > 1 - (1 + \sigma_x^2) \frac{(\log 2)\sigma^2 \max_i \|\beta_{*i}\|_{\Sigma_x}^2}{n\lambda^2 \delta^2}. \end{aligned}$$

Notice that the right hand side doesn't depend on the choice of i and W . Therefore, it holds that with probability at least $1 - (1 + \sigma_x^2) \frac{(\log 2)\sigma^2 \max_i \|\beta_{*i}\|_{\Sigma_x}^2}{n\lambda^2 \delta^2} - 2e^{-cn}$,

$$\frac{\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} < \lambda \left(\epsilon^{I(W, \lambda)} \right)^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*I(W, \lambda)} < \lambda \delta.$$

Finally, changing $\delta \rightarrow \lambda \delta$ gives the result. \square

Lemma 2. *Let x_1, \dots, x_n be i.i.d. centered random vectors in \mathbb{R}^l with covariance matrix Σ , $X = (x_1, \dots, x_n) \in \mathbb{R}^{n \times l}$ and $\hat{\Sigma} = \frac{1}{n} X^\top X$. If x_i is sub-Gaussian with $\|x_i\|_{\psi_2} \leq \sigma_x$, then for any $\beta_1, \dots, \beta_m \in \mathbb{R}^l$ and $\lambda > 0$, there exist constants c, C which only depend on σ_x such that it holds with probability at least $1 - 2e^{-cn}$,*

$$\left(\frac{3}{2} + C \sqrt{\frac{l}{n}} \right)^{-1} \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top (\Sigma + \lambda I_d)^{-1} \Sigma \beta_{*i}$$

$$\begin{aligned}
&\leq \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top (\hat{\Sigma} + \lambda I_d)^{-1} \hat{\Sigma} \beta_{*i} \\
&\leq \left(\frac{3}{2} + C\sqrt{\frac{l}{n}} \right) \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top (\Sigma + \lambda I_d)^{-1} \Sigma \beta_{*i}
\end{aligned}$$

Proof. For a positive semidefinite matrix $A \in \mathbb{R}^{l \times l}$, we define

$$\begin{aligned}
R_i(\beta, A) &:= \|\beta_i - \beta\|_A^2 + \lambda \|\beta\|^2, \\
\hat{\beta}_i(A) &:= (A + \lambda I_l)^{-1} A \beta_i.
\end{aligned}$$

By Lemma 12, there exist constants c, C which only depends on σ_x such that with probability at least $1 - 2e^{-cn}$ for any $i = 1, \dots, m$, it holds

$$\begin{aligned}
&\beta_{*i}^\top (\Sigma + \lambda I_d)^{-1} \Sigma \beta_{*i} \\
&= \min_{\beta \in \mathbb{R}^l} R_i(\beta, \Sigma) \\
&= R_i(\hat{\beta}_i(\Sigma), \Sigma) \\
&\leq R_i(\hat{\beta}_i(\hat{\Sigma}), \Sigma) \\
&= \lambda \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-1} \hat{\Sigma} \beta_i + \lambda^2 \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-1} \Sigma (\hat{\Sigma} + \lambda I_l)^{-1} \beta_i \\
&\quad - \lambda^2 \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-2} \hat{\Sigma} \beta_i \\
&= R_i(\hat{\beta}_i(\hat{\Sigma}), \hat{\Sigma}) + \lambda^2 \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-1} (\Sigma - \hat{\Sigma}) (\hat{\Sigma} + \lambda I_l)^{-1} \beta_i \\
&= \min_{\beta \in \mathbb{R}^l} R_i(\beta, \hat{\Sigma}) + \lambda^2 \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-1} (\Sigma - \hat{\Sigma}) (\hat{\Sigma} + \lambda I_l)^{-1} \beta_i \\
&\leq \min_{\beta \in \mathbb{R}^l} R_i(\beta, \hat{\Sigma}) + \left(\frac{1}{2} + C\sqrt{\frac{l}{n}} \right) \lambda^2 \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-1} \hat{\Sigma} (\hat{\Sigma} + \lambda I_l)^{-1} \beta_i.
\end{aligned}$$

Applying the Sherman-Morrison-Woodbury formula yields

$$\begin{aligned}
(\hat{\Sigma} + \lambda I_l)^{-1} \hat{\Sigma}^{\frac{1}{2}} &= \lambda^{-1} \hat{\Sigma}^{\frac{1}{2}} - \lambda^{-2} \hat{\Sigma}^{\frac{1}{2}} (\lambda^{-1} \hat{\Sigma} + I_l)^{-1} \hat{\Sigma} \\
&= \lambda^{-1} \hat{\Sigma}^{\frac{1}{2}} - \lambda^{-1} \hat{\Sigma}^{\frac{1}{2}} (\lambda^{-1} \hat{\Sigma} + I_l)^{-1} (\lambda^{-1} \hat{\Sigma} + I_l - I_l) \\
&= \hat{\Sigma}^{\frac{1}{2}} (\hat{\Sigma} + \lambda I_l)^{-1}.
\end{aligned}$$

Therefore, it holds

$$\begin{aligned}
\min_{\beta \in \mathbb{R}^l} R_i(\beta, \Sigma) &\leq \min_{\beta \in \mathbb{R}^l} R_i(\beta, \hat{\Sigma}) + \left(\frac{1}{2} + C\sqrt{\frac{l}{n}} \right) \lambda^2 \beta_i^\top (\hat{\Sigma} + \lambda I_l)^{-1} \hat{\Sigma} (\hat{\Sigma} + \lambda I_l)^{-1} \beta_i \\
&= \min_{\beta \in \mathbb{R}^l} R_i(\beta, \hat{\Sigma}) + \left(\frac{1}{2} + C\sqrt{\frac{l}{n}} \right) \lambda^2 \beta_i^\top \hat{\Sigma}^{\frac{1}{2}} (\hat{\Sigma} + \lambda I_l)^{-2} \hat{\Sigma}^{\frac{1}{2}} \beta_i \\
&\leq \min_{\beta \in \mathbb{R}^l} R_i(\beta, \hat{\Sigma}) + \left(\frac{1}{2} + C\sqrt{\frac{l}{n}} \right) \lambda \beta_i^\top \hat{\Sigma}^{\frac{1}{2}} (\hat{\Sigma} + \lambda I_l)^{-1} \hat{\Sigma}^{\frac{1}{2}} \beta_i \\
&= \left(\frac{3}{2} + C\sqrt{\frac{l}{n}} \right) \min_{\beta \in \mathbb{R}^l} R_i(\beta, \hat{\Sigma}).
\end{aligned}$$

Taking the average over $i = 1, \dots, m$ gives

$$\frac{1}{m} \sum_{i=1}^m \min_{\beta \in \mathbb{R}^l} \|\beta_i - \beta\|_\Sigma^2 + \lambda \|\beta\|^2 \leq \left(\frac{3}{2} + C\sqrt{\frac{l}{n}} \right) \frac{1}{m} \sum_{i=1}^m \min_{\beta \in \mathbb{R}^l} \|\beta_i - \beta\|_{\hat{\Sigma}}^2 + \lambda \|\beta\|^2.$$

In the same way as above discussion, we obtain

$$\frac{1}{m} \sum_{i=1}^m \min_{\beta \in \mathbb{R}^l} \|\beta_i - \beta\|_{\hat{\Sigma}}^2 + \lambda \|\beta\|^2 \leq \left(\frac{3}{2} + C \sqrt{\frac{l}{n}} \right) \frac{1}{m} \sum_{i=1}^m \min_{\beta \in \mathbb{R}^l} \|\beta_i - \beta\|_{\Sigma}^2 + \lambda \|\beta\|^2.$$

Combining these two inequalities gives the result. \square

Lemma 3. *Let x_1, \dots, x_n be i.i.d. centered random vectors in \mathbb{R}^d with covariance matrix Σ , $X = (x_1, \dots, x_n) \in \mathbb{R}^{n \times m}$ and $\hat{\Sigma} = \frac{1}{n} X^\top X$. If x_i is sub-Gaussian with $\|x_i\|_{\psi_2} \leq \sigma_x$, $\mu_k(\hat{\Sigma}_{k,k}) \geq \lambda$ and $\mu_1(\hat{\Sigma}_{-k,-k}) \leq \lambda$, then there exists constant c which only depend on σ_x such that it holds with probability at least $1 - 4e^{-cn}$,*

$$\lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k}) \lesssim \lambda \lesssim \mu_k(\Sigma_{k,k}).$$

Proof. Since $\mu_1(\hat{\Sigma}_{-k,-k}) \lesssim \lambda$, from Lemma 11 there exists constant c_1 which only depend on σ_x , such that it holds with probability at least $1 - 2e^{-c_1 n}$,

$$\begin{aligned} \lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k}) &\leq \lambda + \frac{1}{n} (1 + \sigma_x^2) \text{Tr}(\hat{\Sigma}_{-k,-k}) \\ &\leq \lambda + \frac{1}{n} (n - k) (1 + \sigma_x^2) \lambda \\ &\lesssim \lambda. \end{aligned}$$

Since $\mu_k(\hat{\Sigma}_{k,k}) \geq \lambda$, by Lemma 12 there exists constants c_2, C which only depend on σ_x , such that it holds that with probability at least $1 - 2e^{-c_2 n}$,

$$\mu_k(\Sigma_{k,k}) \gtrsim \left(\frac{3}{2} + C \sqrt{\frac{k}{n}} \right) \mu_k(\hat{\Sigma}_{k,k}) \geq \mu_k(\Sigma_{k,k}) \geq \lambda.$$

Combining these two inequalities, we obtain

$$\lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k}) \lesssim \lambda \lesssim \mu_k(\Sigma_{k,k}).$$

\square

B PROOF OF THEOREM 2

Proof. From Lemma 4, for $W = \text{diag}(\mathbf{w}) V^\top$ (V is eigenvectors of Σ_β) such that $(\mathbf{w})_i = \frac{\sqrt{n\lambda\sigma_i}}{\sigma'}$, there exists a constant c which only depends on σ_x such that, with probability at least $1 - 2e^{-cn}$ it holds,

$$\begin{aligned} &\frac{1}{m} \sum_{i=1}^m \min_{\beta \in \mathbb{R}^l} \frac{1}{n} \|\mathbf{y}^{(i)} - XW^\top \beta\|^2 + \lambda \|\beta\|^2 + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\ &= \frac{\lambda}{m} \sum_{i=1}^m \mathbf{y}^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} \mathbf{y}^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(\hat{\Sigma}_{wx} \left(\hat{\Sigma}_{wx} + \lambda I_d \right)^{-1} \right) \\ &= \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \\ &\quad + \frac{\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(X_w (X_w X_w^\top + \lambda I_n)^{-1} X_w^\top \right) \\ &\leq \frac{\lambda}{m} \sum_{i=1}^m \beta_{*i}^\top X^\top (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} + \frac{2\lambda}{m} \sum_{i=1}^m \epsilon^{(i)\top} (X_w X_w^\top + n\lambda I_n)^{-1} X \beta_{*i} \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{nm} \sum_{i=1}^m \epsilon^{(i)\top} \epsilon^{(i)} + \frac{\sigma'^2}{n} \text{Tr} \left(X_w (X_w X_w^\top + \lambda I_n)^{-1} X_w^\top \right) \\
\lesssim & \frac{\lambda}{m} \sum_{i=1}^m \beta_{wi}^\top X_w^\top (X_w X_w^\top + n\lambda I_n)^{-1} X_w \beta_{wi} + \frac{\sigma'^2}{n} \text{Tr} \left(X_w^\top (X_w X_w^\top + \lambda I_n)^{-1} X_w \right) + \sigma^2 \\
= & \frac{2\sigma'^2}{n} \text{Tr} \left(X_w^\top (X_w X_w^\top + n\lambda I_n)^{-1} X_w \right) + \sigma^2. \tag{B.1}
\end{aligned}$$

We write the eigenvectors of $W\Sigma_x W^\top$ corresponding to $\mu_i(W\Sigma_x W^\top)$ as $u_{\beta,i}$ and define $U_\beta := (u_{\beta,1}, \dots, u_{\beta,d})$ and $\Lambda_{x\beta} := \text{diag}(\mu_1(W\Sigma_x W^\top), \dots, \mu_d(W\Sigma_x W^\top))$. Then, we can evaluate (B.1) as follows:

$$\begin{aligned}
& \text{Tr} \left(X_w^\top (X_w X_w^\top + n\lambda I_n)^{-1} X_w \right) \\
= & \text{Tr} \left((X_w U_\beta)^\top \left(X_w U_\beta (X_w U_\beta)^\top + n\lambda I_n \right)^{-1} X_w U_\beta \right) \\
\lesssim & \text{Tr} \left((X_w U_\beta)_k^\top \left(X_w U_\beta (X_w U_\beta)^\top + n\lambda I_n \right)^{-1} (X_w U_\beta)_k \right) \\
& + \text{Tr} \left((X_w U_\beta)_{-k}^\top \left(X_w U_\beta (X_w U_\beta)^\top + n\lambda I_n \right)^{-1} (X_w U_\beta)_{-k} \right).
\end{aligned}$$

For the first term, by Lemma 5,

$$\begin{aligned}
& \text{Tr} \left((X_w U_\beta)_k^\top \left(X_w U_\beta (X_w U_\beta)^\top + n\lambda I_n \right)^{-1} (X_w U_\beta)_k \right) \\
\leq & 2\text{Tr} \left((X_w U_\beta)_k^\top \left((X_w U_\beta)_k (X_w U_\beta)_k^\top + n\lambda I_n \right)^{-1} (X_w U_\beta)_k \right).
\end{aligned}$$

For the second term, we have

$$\begin{aligned}
& \text{Tr} \left((X_w U_\beta)_{-k}^\top \left(X_w U_\beta (X_w U_\beta)^\top + n\lambda I_n \right)^{-1} (X_w U_\beta)_{-k} \right) \\
\leq & \frac{1}{n\lambda} \text{Tr} \left((X_w U_\beta)_{-k}^\top (X_w U_\beta)_{-k} \right).
\end{aligned}$$

From Lemma 11 it holds that with probability at least $1 - 2e^{-c_2 n}$,

$$\text{Tr} \left((X_w U_\beta)_{-k} (X_w U_\beta)_{-k}^\top \right) \leq (1 + \sigma_x^2) \frac{n\lambda}{\sigma'^2} \sum_{i=k+1}^d \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right).$$

Furthermore,

$$\text{Tr} \left((X_w U_\beta)_k^\top \left((X_w U_\beta)_k (X_w U_\beta)_k^\top + n\lambda I_n \right)^{-1} (X_w U_\beta)_k \right) \leq k.$$

Combining these bounds yields

$$\begin{aligned}
& \frac{2\sigma'^2}{n} \text{Tr} \left(X_w^\top (X_w X_w^\top + n\lambda I_n)^{-1} X_w \right) \\
\lesssim & \frac{4k\sigma'^2}{n} + 2(1 + \sigma_x^2) \sum_{i=k+1}^d \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \\
\lesssim & \frac{k\sigma'^2}{n} + \sum_{i=k+1}^d \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right). \tag{B.2}
\end{aligned}$$

Selecting k in (B.2) such that $\mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \leq \frac{\sigma'^2}{n} \leq \mu_k \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)$ gives the result. \square

In the later part of this section, we show lemmas required for Theorem 2.

Lemma 4. Suppose $\epsilon^{(i)}$ are centered vectors whose components are independent and have sub-Gaussian norm at most σ . Then there exists a constant c which only depends on σ such that it holds with probability at least $1 - 2e^{-ct}$

$$\frac{1}{nm} \sum_{i=1}^m \left\| \epsilon^{(i)} \right\|^2 \leq \left(1 + \frac{t}{n} \sigma^2 \right) (\log 2) \sigma^2.$$

Proof. Let define I as $I := \operatorname{argmax}_i \frac{1}{n} \left\| \epsilon^{(i)} \right\|^2$. From Lemma 11, there exists a constant c which only depends on σ such that it holds with probability at least $1 - 2e^{-ct}$

$$\frac{1}{nm} \sum_{i=1}^m \left\| \epsilon^{(i)} \right\|^2 \leq \frac{1}{n} \left\| \epsilon^{(I)} \right\|^2 \leq \left(1 + \frac{t}{n} \sigma^2 \right) \operatorname{Var} \left[\epsilon^{(I)} \right] \leq \left(1 + \frac{t}{n} \sigma^2 \right) (\log 2) \sigma^2.$$

□

Lemma 5. For any matrix $X \in \mathbb{R}^{n \times d}$, it holds that

$$\operatorname{Tr} \left(X_k \left(X X^\top + n \lambda I_n \right) X_k^\top \right) \leq 2 \operatorname{Tr} \left(X_k \left(X_k X_k^\top + n \lambda I_n \right) X_k^\top \right).$$

Proof. We define $\bar{A}_k(\lambda, X) := X_k X_k^\top + n \lambda I_n$. First we show the identity below:

$$X_k^\top \bar{A}_d(\lambda, X)^{-1} + X_k^\top \bar{A}_k(\lambda, X)^{-1} X_{-k} X_{-k}^\top \bar{A}_d(\lambda, X)^{-1} = X_k^\top \bar{A}_k(\lambda, X)^{-1}.$$

Indeed,

$$\begin{aligned} & X_k^\top \bar{A}_d(\lambda, X)^{-1} + X_k^\top \bar{A}_k(\lambda, X)^{-1} X_{-k} X_{-k}^\top \bar{A}_d(\lambda, X)^{-1} \\ &= X_k^\top \left(X_{-k} X_{-k}^\top + \bar{A}_k(\lambda, X) \right)^{-1} + X_k^\top \bar{A}_k(\lambda, X)^{-1} X_{-k} X_{-k}^\top \left(X_{-k} X_{-k}^\top + \bar{A}_k(\lambda, X) \right)^{-1} \\ &= X_k^\top \bar{A}_k(\lambda, X)^{-1} \left(X_{-k} X_{-k}^\top + \bar{A}_k(\lambda, X) \right) \left(X_{-k} X_{-k}^\top + \bar{A}_k(\lambda, X) \right)^{-1} \\ &= X_k^\top \bar{A}_k(\lambda, X)^{-1}. \end{aligned}$$

Thus, multiplying the identity by X_k from the right, we obtain

$$\begin{aligned} \operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_k \right) &= \operatorname{Tr} \left(X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k \right) \\ &\quad + \operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_{-k} X_{-k}^\top \bar{A}_d(\lambda, X)^{-1} X_k \right). \end{aligned}$$

Since, $X_k X_k^\top \bar{A}_k(\lambda, X)^{-1}$ and $X X^\top \bar{A}_d(\lambda, X)^{-1}$ are positive semidefinite, we have

$$\begin{aligned} & \operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_{-k} X_{-k}^\top \bar{A}_d(\lambda, X)^{-1} X_k \right) \\ &= \operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X X^\top \bar{A}_d(\lambda, X)^{-1} X_k \right) \\ &\quad - \operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_k X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k \right) \\ &\geq -\operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_k X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k \right). \end{aligned}$$

$X_k^\top \bar{A}_k(\lambda, X)^{-1} X_k$ and $X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k$ are positive semidefinite. Therefore, it holds

$$\begin{aligned} & -\operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_k X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k \right) \\ &\geq -\mu_1 \left(X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k \right) \operatorname{Tr} \left(X_k^\top \bar{A}_k(\lambda, X)^{-1} X_k \right). \end{aligned}$$

Since $\bar{A}_d(\lambda, X)^{-\frac{1}{2}} X_{-k} X_{-k}^\top \bar{A}_d(\lambda, X)^{-\frac{1}{2}}$ is positive semidefinite, we obtain

$$\begin{aligned} 1 &\geq \mu_1 \left(\bar{A}_d(\lambda, X)^{-\frac{1}{2}} X X^\top \bar{A}_d(\lambda, X)^{-\frac{1}{2}} \right) \\ &= \mu_1 \left(\bar{A}_d(\lambda, X)^{-\frac{1}{2}} X_k X_k^\top \bar{A}_d(\lambda, X)^{-\frac{1}{2}} + \bar{A}_d(\lambda, X)^{-\frac{1}{2}} X_{-k} X_{-k}^\top \bar{A}_d(\lambda, X)^{-\frac{1}{2}} \right) \\ &\geq \mu_1 \left(\bar{A}_d(\lambda, X)^{-\frac{1}{2}} X_k X_k^\top \bar{A}_d(\lambda, X)^{-\frac{1}{2}} \right) \\ &= \mu_1 \left(X_k^\top \bar{A}_d(\lambda, X)^{-1} X_k \right). \end{aligned}$$

Combining these gives the result. □

C PROOF OF PROPOSITION 1

Proof. From Bayes' theorem, we obtain

$$\begin{aligned}
& \mathbb{E}_{\beta_* \sim \mathcal{N}(0, \Sigma_\beta)} \left[\mathbb{E}_{Y \sim \mathcal{N}(X\beta_*, \sigma^2 I)} \left[\|\beta_* - \hat{\beta}\|_{\Sigma_x}^2 \right] \right] \\
&= \int \|\beta_* - \hat{\beta}\|_{\Sigma_x}^2 p(\beta_*) p(y|\beta_*, X) d\beta_* dy \\
&= \int \|\beta_* - \hat{\beta}\|_{\Sigma_x}^2 p(\beta_*|X, y) p(y|X) d\beta_* dy.
\end{aligned}$$

Therefore, it holds

$$\begin{aligned}
& \operatorname{argmin}_{\beta: \text{estimator}} \mathbb{E}_{\beta_* \sim \mathcal{N}(0, \Sigma_\beta)} \left[\mathbb{E}_{Y \sim \mathcal{N}(X\beta_*, \sigma^2 I)} \left[\|\beta_* - \hat{\beta}\|_{\Sigma_x}^2 \right] \right] \\
&= \operatorname{argmin}_{\beta: \text{estimator}} \int \|\beta - \beta_*\|_{\Sigma_x}^2 p(\beta_*|X, y) d\beta_* \\
&= \operatorname{argmin}_{\beta: \text{estimator}} \int (\beta^\top \Sigma_x \beta - 2\beta^\top \Sigma_x \beta_* + \beta_*^\top \Sigma_x \beta_*) p(\beta_*|X, y) d\beta_* \\
&= \operatorname{argmin}_{\beta: \text{estimator}} \left(\beta - \int \beta_* p(\beta_*|X, y) d\beta_* \right)^\top \Sigma_x \left(\beta - \int \beta_* p(\beta_*|X, y) d\beta_* \right) \\
&= \int \beta_* p(\beta_*|X, y) d\beta_*.
\end{aligned}$$

Since $\beta_* \sim \mathcal{N}(0, \Sigma_\beta)$, $\epsilon \sim \mathcal{N}(0, \sigma^2)$ by assumption, it holds

$$\begin{aligned}
& p(\beta_*|x, y) \\
&\propto \exp\left(-\frac{1}{2\sigma^2} \|y - X\beta_*\|^2\right) \exp\left(-\frac{1}{2}\beta_*^\top \Sigma_\beta^{-1} \beta_*\right) \\
&= \exp\left(-\frac{1}{2\sigma^2} \left(\|y\|^2 - 2y^\top X\beta_* + \beta_*^\top (X^\top X + \sigma^2 \Sigma_\beta^{-1}) \beta_*\right)\right) \\
&\propto \exp\left(\frac{1}{2\tilde{\sigma}^2} \left\| \beta_* - (X^\top X + \sigma^2 \Sigma_\beta^{-1})^{-1} X^\top y \right\|_{X^\top X + \sigma^2 \Sigma_\beta^{-1}}^2\right).
\end{aligned}$$

In conclusion, we obtain

$$\int \beta_* p(\beta_*|x, y) d\beta_* = (X^\top X + \sigma^2 \Sigma_\beta^{-1})^{-1} X^\top y.$$

□

D PROOF OF THEOREM 3

Proof. For training data $(x_i, y_i)_{i=1}^n \in \mathbb{R}^d \times \mathbb{R}$, $X = (x_1, \dots, x_n)^\top$ and $y(\beta_*, \epsilon) = (y_1, \dots, y_n)^\top$ such that $y(\beta_*, \epsilon) = X\beta_* + \epsilon$, we define $\hat{\beta}(W, \beta_*, \epsilon) = W^\top (WX^\top XW^\top + n\lambda I_d)^{-1} WX^\top y(\beta_*, \epsilon)$. It holds that

$$\begin{aligned}
& \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{x, \epsilon_i} \left[\left(x^\top \beta_{*i} - x^\top W^\top \hat{\beta}_i(W) \right)^2 \right] \\
&= \mathbb{E}_{x \sim \mathcal{N}(0, \Sigma_x), \epsilon \sim \mathcal{N}(0, \tilde{\sigma}^2), \beta_* \sim \mathcal{N}(0, \Sigma_\beta)} \left[\left(x^\top \beta_* - x^\top W^\top \hat{\beta}(W, \beta_*, \epsilon) \right)^2 \right] \\
&\geq \min_{\beta: \text{estimator}} \mathbb{E}_{x \sim \mathcal{N}(0, \Sigma_x), \epsilon \sim \mathcal{N}(0, \tilde{\sigma}^2), \beta_* \sim \mathcal{N}(0, \Sigma_\beta)} \left[\left(x^\top \beta_* - x^\top \beta \right)^2 \right] \\
&= R_{\text{opt}}(X, \tilde{\sigma}).
\end{aligned} \tag{D.1}$$

From Proposition 1, for $W_B = \frac{\sqrt{n\lambda}}{\tilde{\sigma}} \Sigma_\beta^{\frac{1}{2}}$, we can write Bayes estimator as

$$\begin{aligned}\beta_{\text{Bayes}}(\beta_*, \epsilon) &= \left(X^\top X + \tilde{\sigma}^2 \Sigma_\beta^{-1} \right)^{-1} X^\top y(\beta_*, \epsilon) \\ &= W_B^\top \left(W_B X^\top X W_B^\top + n\lambda I_d \right)^{-1} W_B X^\top y(\beta_*, \epsilon).\end{aligned}$$

Therefore for $\Sigma_{x,B} = \frac{n\lambda}{\tilde{\sigma}^2} \Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}}$ and $\hat{\Sigma}_{x,B} = \frac{n\lambda}{\tilde{\sigma}^2} \Sigma_\beta^{\frac{1}{2}} \hat{\Sigma}_x \Sigma_\beta^{\frac{1}{2}}$, we can write the right hand side of (D.1) as

$$\begin{aligned}R_{\text{opt}}(X, \tilde{\sigma}) &= \mathbb{E}_{x \sim \mathcal{N}(0, \Sigma_x), \epsilon \sim \mathcal{N}(0, \tilde{\sigma}^2), \beta_* \sim \mathcal{N}(0, \Sigma_\beta)} \left[\left(x^\top \beta_* - x^\top \beta_{\text{Bayes}}(\beta_*, \epsilon) \right)^2 \right] \\ &= B_{\text{Bayes}} + V_{\text{Bayes}},\end{aligned}$$

where,

$$\begin{aligned}B_{\text{Bayes}} &:= \text{Tr} \left(\left(\left(n\hat{\Sigma}_{x,B} + n\lambda I_d \right)^{-1} n\hat{\Sigma}_{x,B} - I_d \right) \Sigma_{x,B} \right), \\ V_{\text{Bayes}} &= \mathbb{E}_\epsilon \left[\left\| W_B X^\top \left(X W_B^\top W_B X^\top + n\lambda I_n \right)^{-1} \epsilon \right\|_{\Sigma_{x,B}}^2 \right].\end{aligned}$$

For k in Assumption 3, we define ρ_k such that $\frac{n\lambda}{\tilde{\sigma}^2} \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \rho_k = \left(\lambda + \frac{\lambda}{\tilde{\sigma}^2} \sum_{i=k+1}^d \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \right)$. For the variance term, from Assumption 2 and Lemma 7, there exists some absolute constant c such that for any t which satisfies $c < t < n$ and $k + 2\sigma^2 + \sqrt{kt}\sigma^2 < n/2$, it holds that with probability at least $1 - 20e^{-t/c}$,

$$\begin{aligned}V_{\text{Bayes}} &= \mathbb{E}_\epsilon \left[\left\| W_B X^\top \left(X W_B^\top W_B X^\top + n\lambda I_n \right)^{-1} \epsilon \right\|_{\Sigma_{x,B}}^2 \right] \\ &\geq \frac{\tilde{\sigma}^2}{cn} \sum_{i=1}^d \min \left\{ 1, \frac{\frac{n^2 \lambda^2}{\tilde{\sigma}^4} \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2}{\sigma_x^4 \frac{n^2 \lambda^2}{\tilde{\sigma}^4} \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2 (\rho_k + 2)^2} \right\} \\ &\gtrsim \frac{\tilde{\sigma}^2}{cn} \sum_{i=1}^d \min \left\{ 1, \frac{\frac{n^2 \lambda^2}{\tilde{\sigma}^4} \mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2}{\sigma_x^4 \frac{n^2 \lambda^2}{\tilde{\sigma}^4} \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2 \rho_k^2} \right\} \\ &= \frac{\tilde{\sigma}^2}{cn} \sum_{i=1}^d \min \left\{ 1, \frac{\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2}{\sigma_x^4 \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2 \rho_k^2} \right\} \\ &\geq \frac{\tilde{\sigma}^2}{cn} \sum_{i=1}^d \frac{\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2}{\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2 + \sigma_x^4 \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2 \rho_k^2} \\ &\gtrsim \frac{\tilde{\sigma}^2}{cn} \sum_{i=1}^d \frac{\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2}{\left(\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) + \sigma_x^2 \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \rho_k \right)^2} \\ &\gtrsim \frac{\tilde{\sigma}^2}{cn} \sum_{i=1}^d \frac{\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right)^2}{\left(\mu_i \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) + \mu_{k+1} \left(\Sigma_\beta^{\frac{1}{2}} \Sigma_x \Sigma_\beta^{\frac{1}{2}} \right) \rho_k \right)^2}.\end{aligned}$$

In the last inequality, we use Assumption 2. For the bias term, we define eigenvectors $u_{B,1}, \dots, u_{B,d}$ corresponding to $\mu_1(W_B \Sigma_x W_B^\top), \dots, \mu_1(W_B \Sigma_x W_B^\top)$ and $U_B = (u_{B,1}, \dots, u_{B,d})$. First, for any $j > K$, we need to evaluate the smallest eigenvalue of $A_{-j}(X W_B^\top U_B, \lambda) =$

$(XW_B^\top U_B)_{j-1}(XW_B^\top U_B)_{j-1}^\top + (XW_B^\top U_B)_{-j}(XW_B^\top U_B)_{-j}^\top + n\lambda I_n$ in order to apply Proposition 6. From Assumption 3, we have

$$n\lambda + \frac{n\lambda}{\tilde{\sigma}^2} \sum_{i=k+1}^d \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \leq 2n\lambda \leq 2\mu_n \left(A_{-j} \left(XW_B^\top U_B, \lambda \right) \right).$$

Therefore, since $W_B^{-\top} \Sigma_{\beta} W_B^{-1} = \frac{\tilde{\sigma}^2}{n\lambda} I_d$, from Lemma 6 there exists some absolute constant c' such that for any t which satisfies $t < \frac{n}{2\tilde{\sigma}^2}$, it holds that with probability at least $1 - 4e^{-t/c'}$,

$$\begin{aligned} B_{\text{Bayes}} &:= \text{Tr} \left(\left(\left(n\hat{\Sigma}_{x,B} + n\lambda I_d \right)^{-1} n\hat{\Sigma}_{x,B} - I_d \right)^2 \Sigma_{x,B} \right) \\ &\geq \frac{1}{2} \sum_{i=1}^d \frac{\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{\left(1 + \frac{\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{2\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k} \right)^2} \\ &\gtrsim \sum_{i=1}^d \frac{\left(\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k \right)^2 \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{\left(\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k + \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right)^2} \\ &\gtrsim \sum_{i=1}^d \frac{\tilde{\sigma}^2 \left(\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k \right)^2 \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{\left(\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k + \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right)^2}. \end{aligned}$$

In the last inequality, we use Assumption 2. Combining these two bounds yields

$$\begin{aligned} &B_{\text{Bayes}} + V_{\text{Bayes}} \\ &\gtrsim \sum_{i=1}^d \frac{\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \tilde{\sigma}^2 n \left(\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k \right)^2 + \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{n \left(\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k + \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right)^2} \\ &= \sum_{i=1}^d \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \tilde{\sigma}^2 \frac{\left(n\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k \right)^2 + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{\left(n\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right)^2}. \quad (\text{D.2}) \end{aligned}$$

By assumption 2 and 3, $n\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k = \tilde{\sigma}^2 + \sum_{i=k+1}^d \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \leq 2\tilde{\sigma}^2$, and (D.2) is a decreasing function as long as $n\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k \leq 2\tilde{\sigma}^2 \leq 1$. Therefore, we have

$$\begin{aligned} B_{\text{Bayes}} + V_{\text{Bayes}} &\gtrsim \sum_{i=1}^d \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \tilde{\sigma}^2 \frac{\left(n\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k \right)^2 + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{\left(n\mu_{k+1} \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \rho_k + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right)^2} \\ &\geq \sum_{i=1}^d \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \tilde{\sigma}^2 \frac{2\tilde{\sigma}^2 + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)}{\left(2\tilde{\sigma}^2 + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right)^2} \\ &= \sum_{i=1}^d \frac{\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \tilde{\sigma}^2}{2\tilde{\sigma}^2 + n\mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right)} \\ &\gtrsim \sum_{i=1}^d \min \left\{ \frac{\tilde{\sigma}^2}{n}, \mu_i \left(\Sigma_{\beta}^{\frac{1}{2}} \Sigma_x \Sigma_{\beta}^{\frac{1}{2}} \right) \right\}. \end{aligned}$$

□

In the later part of the section we show lemmas needed to prove Theorem 3.

Lemma 6 (Lower bound of bias term). *Suppose that Σ_x and Σ_β have the same eigenvectors, i.e., $U = V$ and that it is known for some k, δ, L that for any $j > k$ with probability at least $1 - \delta$, $\mu_n(A_{-j}(X, \lambda)) \geq \frac{1}{L}(n\lambda + \sum_{i>k} \lambda_i)$. Then for some absolute constant c for any non-negative $t < \frac{n}{2\sigma_x^2}$ it holds that with probability at least $1 - 2\delta - 4e^{-t/c}$,*

$$\begin{aligned} \mathcal{B} &:= \text{Tr} \left(\Sigma_\beta^{\frac{1}{2}} \left((X^\top X + n\lambda I_d)^{-1} X^\top X - I_d \right) \Sigma_x \left((X^\top X + n\lambda I_d)^{-1} X^\top X - I_d \right) \Sigma_\beta^{\frac{1}{2}} \right) \\ &\geq \frac{1}{2} \sum_{i=1}^d \frac{\lambda_i \sigma_i^2}{\left(1 + \frac{\lambda_i}{2L\lambda_{k+1}\rho_k} \right)^2}. \end{aligned}$$

Proof. Let $X_u := XU$, $\Lambda := \mathbb{E}[U^\top x x^\top U] = U^\top \Sigma_x U = \text{diag}(\lambda_1, \dots, \lambda_d)$, $\Sigma = U^\top \Sigma_\beta U = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$. In the same way as Lemma 15 in (Tsigler and Bartlett, 2020) we can write the bias term as

$$\begin{aligned} \mathcal{B} &= \text{Tr} \left(\Sigma_\beta^{\frac{1}{2}} \left(I_d - X^\top (X X^\top + n\lambda I_n)^{-1} X \right) \Sigma_x \left(I_d - X^\top (X X^\top + n\lambda I_n)^{-1} X \right) \Sigma_\beta^{\frac{1}{2}} \right) \\ &= \text{Tr} \left(\Sigma_\beta^{\frac{1}{2}} \left(I_d - X_u^\top (X_u X_u^\top + n\lambda I_n)^{-1} X_u \right) \Lambda \left(I_d - X_u^\top (X_u X_u^\top + n\lambda I_n)^{-1} X_u \right) \Sigma_\beta^{\frac{1}{2}} \right). \\ &= \sum_{i=1}^d \left(\left(I_d - X_u^\top (X_u X_u^\top + n\lambda I_n)^{-1} X_u \right) \Lambda \left(I_d - X_u^\top (X_u X_u^\top + n\lambda I_n)^{-1} X_u \right) \right)_{i,i} \sigma_i^2. \end{aligned}$$

Λ is a diagonal matrix and so the same discussion as Lemma 15 in (Tsigler and Bartlett, 2020) gives the result. \square

Lemma 7 (Lower bound of variance term). *Suppose that components the rows of XU are independent and $\text{Var}[\epsilon^{(i)}] = \sigma^2$. Then for some absolute constant c , for any t, k such that $t > c$ and $k + 2\sigma_x^2 + \sqrt{kt}\sigma_x^2 < n/2$ it holds that with probability at least $1 - 20e^{-t/c}$,*

$$\mathcal{V} := \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\epsilon^{(i)}} \left[\left\| X^\top (X X^\top + \lambda I_n)^{-1} \epsilon^{(i)} \right\|_{\Sigma_x}^2 \right] \geq \frac{\sigma^2}{cn} \sum_{i=1}^d \min \left\{ 1, \frac{\lambda_i^2}{\sigma_x^4 \lambda_{k+1}^2 (\rho_k + 2)^2} \right\}.$$

Proof. The variance term can be written as

$$\begin{aligned} \mathcal{V} &= \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\epsilon^{(i)}} \left[\left\| X^\top (X X^\top + \lambda I_n)^{-1} \epsilon^{(i)} \right\|_{\Sigma_x}^2 \right] \\ &= \mathbb{E}_{\epsilon^{(1)}} \left[\left\| X^\top (X X^\top + \lambda I_n)^{-1} \epsilon^{(1)} \right\|_{\Sigma_x}^2 \right] \\ &= \mathbb{E}_{\epsilon^{(1)}} \left[\left\| (XU)^\top \left((XU)(XU)^\top + \lambda I_n \right)^{-1} \epsilon^{(1)} \right\|_{U^\top \Sigma_x U}^2 \right]. \end{aligned}$$

Since $\mathbb{E}[U^\top x x^\top U] = U^\top \Sigma_x U = \text{diag}(\lambda_1, \dots, \lambda_d)$, the same discussion as Lemma 14 in (Tsigler and Bartlett, 2020) gives the result. \square

E EVALUATION OF BIAS AND VARIANCE OF RIDGE REGRESSION IN MULTIVARIATE LINEAR REGRESSION

In this section, for $\lambda \in \mathbb{R}$ and $\eta \in \mathbb{R}^n$, we write $\hat{\beta}(\lambda, \eta) := X^\top (X^\top + n\lambda I_n)^{-1} \eta$.

Proposition 2. *Suppose the condition number of $A_k(X, \lambda)$ is at most L , then there exists (large) constants c_x , which only depends on σ_x such that for all t which satisfies*

1. $t \in (1, n/c_x)$,
2. $\sqrt{k} + \sqrt{t} \leq \sqrt{n}/c_x$,

if k and λ satisfies $\frac{1}{\mu_1(\Sigma_{-k,-k})} \left(\lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k}) \right) = \Omega\left(\frac{1}{L}\right)$, then it holds that with probability at least $1 - 20e^{-t/c_x}$,

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x^\top \beta_{*i} - x^\top \hat{\beta}(\lambda, y_i) \right)^2 \right] \lesssim B + V, \\ & \frac{B}{L^3 \max \left\{ L, \frac{\text{Tr}(\Sigma_{-k,-k})}{n\lambda + \text{Tr}(\Sigma_{-k,-k})} \right\}} = \text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{-k,-k} \right) \left(\lambda + \frac{1}{n} \text{Tr}((\Sigma_x)_{-k,-k}) \right)^2 \\ & \quad + \text{Tr}((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k}), \\ & \frac{V}{t\sigma^2 L^2} = \frac{k}{n} + \frac{1}{n} \frac{\text{Tr}((\Sigma_x)_{-k,-k}^2)}{\left(\lambda + \frac{1}{n} \text{Tr}((\Sigma_x)_{-k,-k}) \right)^2}. \end{aligned}$$

Alternatively, if $\lambda > 0$ and it is also known that for some $\delta < 1 - 4^{-n/c_x^2}$ with probability at least $1 - \delta$, the condition number of $A_k(X, \lambda)$ is at most L , then it holds that with probability at least $1 - 20e^{-t/c_x} - \delta$,

$$\begin{aligned} & \frac{B}{L^4} = \text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{-k,-k} \right) \left(\lambda + \frac{1}{n} \text{Tr}((\Sigma_x)_{-k,-k}) \right)^2 + \text{Tr}((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k}), \\ & \frac{V}{t\sigma^2 L^2} = \frac{k}{n} + \frac{1}{n} \frac{\text{Tr}((\Sigma_x)_{-k,-k}^2)}{\left(\lambda + \frac{1}{n} \text{Tr}((\Sigma_x)_{-k,-k}) \right)^2}, \\ & \frac{n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})}{n\mu_1((\Sigma_x)_{-k,-k})} \geq \frac{1}{c_x L}. \end{aligned}$$

Proof.

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x^\top \beta_{*i} - x^\top X^\top (XX^\top + n\lambda I_n)^{-1} y_i \right)^2 \right] \\ & \leq \frac{2}{m} \sum_{i=1}^m \left\| \beta_{*i} - \hat{\beta}(\lambda, X\beta_{*i}) \right\|_{\Sigma_x}^2 + \frac{2}{m} \sum_{i=1}^m \left\| X^\top (XX^\top + n\lambda I_n)^{-1} \epsilon^{(i)} \right\|_{\Sigma_x}^2. \end{aligned}$$

The first term corresponds to the bias and the second term corresponds to that of the variance. From Lemma 8 and Lemma 9, for all $t > 1$ it holds that with probability at least $1 - 4e^{-c_1 t}$,

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \left\| \beta_{*i} - \hat{\beta}(\lambda, y_i) \right\|_{\Sigma_x}^2 \\ & \lesssim \frac{\mu_1(A_k(X, \lambda)^{-1})^2}{\mu_n(A_k(X, \lambda)^{-1})^2} \frac{\mu_1((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2})}{\mu_k((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2})^2} \left(\frac{1}{m} \sum_{i=1}^m \|X_{-k}(\beta_{*i})_{-k}\|^2 \right) \\ & \quad + \frac{\text{Tr}((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k})}{\mu_n(A_k(X, \lambda)^{-1})^2 \mu_k((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2})^2} \\ & \quad + \|X_{-k}(\Sigma_x)_{-k,-k} X_{-k}^\top\| \left(\mu_1(A_1(X, \lambda)^{-1})^2 \left(\frac{1}{m} \sum_{i=1}^m \|X_{-k}(\beta_{*i})_{-k}\|^2 \right) \right) \\ & \quad + \frac{\mu_1(A_k(X, \lambda)^{-1})^2}{\mu_n(A_k(X, \lambda)^{-1})^2} \frac{\mu_1((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2})}{\mu_k((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2})^2} \text{Tr}((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k}) \end{aligned}$$

$$\begin{aligned}
& + \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right) + C_1 \sigma^2 t \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{\text{Tr} \left(X_k (\Sigma_x)_{k,k}^{-1} X_k^\top \right)}{\mu_k \left((\Sigma_x)_{k,k}^{-\frac{1}{2}} X_k^\top X_k (\Sigma_x)_{k,k}^{-\frac{1}{2}} \right)^2} \\
& + C_1 \sigma^2 t \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 \text{Tr} \left(X_{-k} (\Sigma_x)_{-k,-k} X_{-k}^\top \right).
\end{aligned}$$

The matrix $X_k (\Sigma_x)_{k,k}^{-1/2} \in \mathbb{R}^{k \times n}$ has n i.i.d. columns with isotropic sub-gaussian distribution in \mathbb{R}^k . By theorem 5.39 in (Vershynin, 2012) there exist constants c'_x, C'_x (which only depends on σ_x) such that for every $t > 0$ s.t. $\sqrt{n} - C'_x \sqrt{k} - \sqrt{t} > 0$, it holds that with probability at least $1 - 2e^{-c'_x t}$,

$$\begin{aligned}
\mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right) & \geq \left(\sqrt{n} - C'_x \sqrt{k} - \sqrt{t} \right)^2, \\
\mu_1 \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right) & \leq \left(\sqrt{n} + C'_x \sqrt{k} + \sqrt{t} \right)^2.
\end{aligned}$$

$\text{Tr} \left(X_k (\Sigma_x)_{k,k}^{-1} X_k^\top \right)$ is the sum of l_2 norms of n i.i.d. isotropic vectors in \mathbb{R}^k , and $\text{Tr} \left(X_{-k} (\Sigma_x)_{-k,-k} X_{-k}^\top \right)$ is the sum of squared norms of n i.i.d. sub-gaussian vectors with covariance $(\Sigma_x)_{-k,-k}^2$. By Lemma 11, it holds that with probability at least $1 - 4e^{-c_2 t}$,

$$\begin{aligned}
\text{Tr} \left(X_k (\Sigma_x)_{k,k}^{-1} X_k^\top \right) & \leq (n + t\sigma_x^2) k, \\
\text{Tr} \left(X_{-k} (\Sigma_x)_{-k,-k}^2 X_{-k}^\top \right) & \leq (n + t\sigma_x^2) \text{Tr} \left((\Sigma_x)_{-k,-k}^2 \right).
\end{aligned}$$

By Lemma 20 in (Tsigler and Bartlett, 2020) it holds that with probability at least $1 - 6e^{-t/c_3}$,

$$\|X_{-k} (\Sigma_x)_{-k,-k} X_{-k}^\top\| \leq c_3 \sigma_x^2 \left((t+n) \mu_1 \left((\Sigma_x)_{-k,-k} \right)^2 + \text{Tr} \left((\Sigma_x)_{-k,-k}^2 \right) \right).$$

$\frac{1}{m} \sum_{i=1}^m \|X_{-k} (\beta_{*i})_{-k}\|^2 = \text{Tr} \left(X_{-k} (\Sigma_\beta)_{-k,-k} X_{-k}^\top \right)$ is the sum of l_2 norms of n i.i.d. sub-gaussian vectors with covariance $(\Sigma_x)_{-k,-k}^{1/2} (\Sigma_\beta)_{-k,-k} (\Sigma_x)_{-k,-k}^{1/2}$. We can apply Lemma 11 to get that for any $t \in (0, n)$ it holds that with probability at least $1 - 4e^{-c_2 t}$,

$$\frac{1}{m} \sum_{i=1}^m \|X_{-k} (\beta_{*i})_{-k}\|^2 = (n + t\sigma_x^2) \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right).$$

Combining the above bounds gives that for some constant C'_x which only depends on σ_x , it holds that

$$\begin{aligned}
\frac{B}{C'_x} & \leq \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right) \left(1 + \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \right) \\
& + \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right) \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 n^2 \mu_1 \left((\Sigma_x)_{-k,-k} \right)^2 \\
& + \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right) \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 n \text{Tr} \left((\Sigma_x)_{-k,-k}^2 \right) \\
& + \frac{\text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right)}{n^2 \mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \\
& + \text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right) \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \mu_1 \left((\Sigma_x)_{-k,-k} \right)^2 \\
& + \text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right) \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{1}{n} \text{Tr} \left((\Sigma_x)_{-k,-k}^2 \right), \tag{E.1}
\end{aligned}$$

$$\frac{V}{C'_x \sigma^2 t} \leq \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 k}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2 n} + n \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 \text{Tr} \left((\Sigma_x)_{-k,-k}^2 \right). \tag{E.2}$$

Note that by Lemma 21 in (Tsigler and Bartlett, 2020) there exists a constant c_4 such that with probability at least $1 - \delta - 2e^{-c_4 t}$,

$$\begin{aligned} \frac{n - t\sigma_x^2}{nL} (n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})) &\leq \mu_k(A_k(X, \lambda)^{-1}) \\ &\leq \mu_1(A_k(X, \lambda)^{-1}) \leq \frac{(n + t\sigma_x^2)L}{n} (n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})). \end{aligned}$$

Plugging that into (E.1) and (E.2) gives

$$\begin{aligned} \frac{B}{C'_x L^2} &\leq \text{Tr}((\Sigma_x)_{-k,-k}(\Sigma_\beta)_{-k,-k}) \left(1 + \frac{n^2 \mu_1((\Sigma_x)_{-k,-k})^2 + n \text{Tr}((\Sigma_x)_{-k,-k}^2)}{(n\lambda + \text{Tr}((\Sigma_x)_{-k,-k}))^2} \right) \\ &\quad + \text{Tr}((\Sigma_x)_{k,k}^{-1}(\Sigma_\beta)_{k,k}) \left(\frac{n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})}{n} \right)^2 \\ &\quad + \text{Tr}((\Sigma_x)_{k,k}^{-1}(\Sigma_\beta)_{k,k}) \left(\mu_1((\Sigma_x)_{-k,-k})^2 + \frac{1}{n} \text{Tr}((\Sigma_x)_{-k,-k}^2) \right), \quad (\text{E.3}) \\ \frac{V}{C''_x \sigma^2 t L^2} &\leq \frac{k}{n} + n \frac{\text{Tr}((\Sigma_x)_{-k,-k}^2)}{(n\lambda + \text{Tr}((\Sigma_x)_{-k,-k}))^2}. \end{aligned}$$

For the first assertion, we have $\frac{1}{\mu_1(\Sigma_{-k,-k})} (\lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k})) = \Omega(\frac{1}{L})$ by assumption. Therefore, it holds

$$\text{Tr}((\Sigma_x)_{-k,-k}^2) \leq \mu_1(\Sigma_{-k,-k}) \text{Tr}((\Sigma_x)_{-k,-k}) \lesssim \frac{L}{n} (n\lambda + \text{Tr}(\Sigma_{-k,-k})) \text{Tr}(\Sigma_{-k,-k}).$$

Plugging this inequality and $\frac{1}{\mu_1(\Sigma_{-k,-k})} (\lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k})) = \Omega(\frac{1}{L})$ into (E.3), it holds that

$$\begin{aligned} \frac{B}{C'_x L^2} &\lesssim \text{Tr}((\Sigma_x)_{-k,-k}(\Sigma_\beta)_{-k,-k}) \left(1 + L^2 + \frac{L \text{Tr}(\Sigma_{-k,-k})}{n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})} \right) \\ &\quad + \text{Tr}((\Sigma_x)_{k,k}^{-1}(\Sigma_\beta)_{k,k}) \left(\frac{n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})}{n} \right)^2 \\ &\quad + \text{Tr}((\Sigma_x)_{k,k}^{-1}(\Sigma_\beta)_{k,k}) \left(L^2 \left(\lambda + \frac{1}{n} \text{Tr}(\Sigma_{-k,-k}) \right)^2 \right. \\ &\quad \left. + \frac{L}{n^2} (n\lambda + \text{Tr}(\Sigma_{-k,-k})) \text{Tr}(\Sigma_{-k,-k}) \right), \end{aligned}$$

and, this gives the first assertion.

For the second assertion, when $\lambda \geq 0$, recall that Lemma 21 in (Tsigler and Bartlett, 2020) also says that $\delta < 1 - 4e^{-c_4 s}$ implies

$$\frac{n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})}{n\mu_1((\Sigma_x)_{-k,-k})} \geq \frac{1}{L} \cdot \frac{1 - s\sigma_x^2/n}{1 + s\sigma_x^2/n}.$$

Moreover, in the same way as the first assertion, it holds

$$\text{Tr}((\Sigma_x)_{-k,-k}^2) \leq \mu_1((\Sigma_x)_{-k,-k}) \text{Tr}((\Sigma_x)_{-k,-k}) \leq \frac{L}{n} (n\lambda + \text{Tr}((\Sigma_x)_{-k,-k}))^2 \frac{1 + s\sigma_x^2/n}{1 - s\sigma_x^2/n}.$$

Plugging it into (E.3) gives the upper bound of the bias in the second assertion. In addition, when $\lambda \geq 0$, it holds that

$$\frac{V}{C''_x \sigma^2 t L^2} \leq \frac{k}{n} + n \frac{\text{Tr}((\Sigma_x)_{-k,-k}^2)}{(n\lambda + \text{Tr}((\Sigma_x)_{-k,-k}))^2} \leq \frac{k}{n} + n \frac{\text{Tr}((\Sigma_x)_{-k,-k})}{n\lambda + \text{Tr}((\Sigma_x)_{-k,-k})},$$

and this gives the second assertion. \square

Lemma 8 (Bias term). *In the same setting as Proposition 2, it holds*

$$\begin{aligned}
& \frac{1}{m} \sum_{i=1}^m \left\| \beta_{*i} - \hat{\beta}(\lambda, X\beta_{*i}) \right\|_{\Sigma_x}^2 \\
\lesssim & \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{\mu_1 \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)}{\mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)^2} \left(\frac{1}{m} \sum_{i=1}^m \|X_{-k}(\beta_{*i})_{-k}\|^2 \right) \\
& + \frac{\text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right)}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2 \mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)^2} \\
& + \|X_{-k}(\Sigma_x)_{-k,-k} X_{-k}^\top\| \left(\mu_1 \left(A_1(X, \lambda)^{-1} \right)^2 \left(\frac{1}{m} \sum_{i=1}^m \|X_{-k}(\beta_{*i})_{-k}\|^2 \right) \right. \\
& \left. + \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{\mu_1 \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)}{\mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)^2} \text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right) \right) \\
& + \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right).
\end{aligned}$$

Proof.

$$\begin{aligned}
& \frac{1}{m} \sum_{i=1}^m \left\| \beta_{*i} - \hat{\beta}(\lambda, X\beta_{*i}) \right\|_{\Sigma_x}^2 \\
= & \frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x^\top \beta_{*i} - x^\top X^\top (XX^\top + n\lambda I_n) X\beta_{*i} \right)^2 \right] \\
= & \frac{1}{m} \sum_{i=1}^m \mathbb{E}_x \left[\left(x_k^\top (\beta_{*i})_k + x_{-k}^\top (\beta_{*i})_{-k} - x_k^\top X_k^\top (XX^\top + n\lambda I_n) X\beta_{*i} \right. \right. \\
& \left. \left. - x_{-k}^\top X_{-k}^\top (XX^\top + n\lambda I_n) X\beta_{*i} \right)^2 \right] \\
\lesssim & \frac{1}{m} \sum_{i=1}^m \left\| (\beta_{*i})_k - X_k^\top (XX^\top + n\lambda I_n) X\beta_{*i} \right\|_{(\Sigma_x)_k}^2 \\
& + \frac{1}{m} \sum_{i=1}^m \left\| X_{-k}^\top (XX^\top + n\lambda I_n) X\beta_{*i} \right\|_{(\Sigma_x)_{-k}}^2 + \frac{1}{m} \sum_{i=1}^m \|(\beta_{*i})_{-k}\|_{(\Sigma_x)_{-k}}^2.
\end{aligned}$$

By the same discussion of chapter F.1 and F.2 in (Tsigler and Bartlett, 2020), we obtain the following bound:

$$\begin{aligned}
& \frac{1}{m} \sum_{i=1}^m \left\| \beta_{*i} - \hat{\beta}(\lambda, X\beta_{*i}) \right\|_{\Sigma_x}^2 \\
\lesssim & \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{\mu_1 \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)}{\mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)^2} \left(\frac{1}{m} \sum_{i=1}^m \|X_{-k}(\beta_{*i})_{-k}\|^2 \right) \\
& + \frac{\text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right)}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2 \mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)^2} \\
& + \|X_{-k}(\Sigma_x)_{-k,-k} X_{-k}^\top\| \left(\mu_1 \left(A_1(X, \lambda)^{-1} \right)^2 \left(\frac{1}{m} \sum_{i=1}^m \|X_{-k}(\beta_{*i})_{-k}\|^2 \right) \right.
\end{aligned}$$

$$\begin{aligned}
& + \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{\mu_1 \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)}{\mu_k \left((\Sigma_x)_{k,k}^{-1/2} X_k^\top X_k (\Sigma_x)_{k,k}^{-1/2} \right)^2} \text{Tr} \left((\Sigma_x)_{k,k}^{-1} (\Sigma_\beta)_{k,k} \right) \\
& + \text{Tr} \left((\Sigma_x)_{-k,-k} (\Sigma_\beta)_{-k,-k} \right).
\end{aligned}$$

□

Lemma 9 (Variance term). *In the same setting as Proposition 2, for all $t > 1$ there exist constants c, C such that with probability at least $1 - 4e^{-ct}$,*

$$\begin{aligned}
& \frac{1}{m} \sum_{i=1}^m \left\| X^\top (X X^\top + n \lambda I_n)^{-1} \epsilon^{(i)} \right\| \\
& \leq C \sigma^2 t \frac{\mu_1 \left(A_k(X, \lambda)^{-1} \right)^2}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2} \frac{\text{Tr} \left(X_k (\Sigma_x)_{k,k}^{-1} X_k^\top \right)}{\mu_k \left((\Sigma_x)_{k,k}^{-\frac{1}{2}} X_k^\top X_k (\Sigma_x)_{k,k}^{-\frac{1}{2}} \right)^2} \\
& \quad + C \sigma^2 t \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 \text{Tr} \left(X_{-k} (\Sigma_x)_{-k,-k} X_{-k}^\top \right).
\end{aligned}$$

Proof. From Lemma 12 in (Tsigler and Bartlett, 2020), we obtain

$$\begin{aligned}
V & = \frac{1}{m} \sum_{i=1}^m \left\| X^\top (X X^\top + n \lambda I_n)^{-1} \epsilon^{(i)} \right\|_{\Sigma_x}^2 \\
& \leq \frac{1}{m} \sum_{i=1}^m \frac{\epsilon^{(i)\top} A_k(X, \lambda)^{-1} X_k (\Sigma_x)_{k,k}^{-1} X_k^\top A_k(X, \lambda)^{-1} \epsilon^{(i)}}{\mu_n \left(A_k(X, \lambda)^{-1} \right)^2 \mu_k \left((\Sigma_x)_{k,k}^{-\frac{1}{2}} X_k^\top X_k (\Sigma_x)_{k,k}^{-\frac{1}{2}} \right)^2} \\
& \quad + \frac{1}{m} \sum_{i=1}^m \epsilon^{(i)\top} A_0(X, \lambda)^{-1} X_{-k} \Sigma_{-k} X_{-k}^\top A_0(X, \lambda)^{-1} \epsilon^{(i)}. \tag{E.4}
\end{aligned}$$

From Lemma 13, for all $t > 1$ it holds that with probability at least $1 - 4e^{-ct}$,

$$\begin{aligned}
& \frac{1}{m} \sum_{i=1}^m \epsilon^{(i)\top} A_k(X, \lambda)^{-1} X_k (\Sigma_x)_{k,k}^{-1} X_k^\top A_k(X, \lambda)^{-1} \epsilon^{(i)} \\
& \leq C \sigma^2 t \text{Tr} \left(A_k(X, \lambda)^{-1} X_k (\Sigma_x)_{k,k}^{-1} X_k^\top A_k(X, \lambda)^{-1} \right) \\
& \leq C \sigma^2 t \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 \text{Tr} \left(X_k (\Sigma_x)_{k,k}^{-1} X_k^\top \right), \\
& \frac{1}{m} \sum_{i=1}^m \epsilon^{(i)\top} A_0(X, \lambda)^{-1} X_{-k} (\Sigma_x)_{-k,-k}^{-1} X_{-k}^\top A_0(X, \lambda)^{-1} \epsilon^{(i)} \\
& \leq C \sigma^2 t \text{Tr} \left(A_0(X, \lambda)^{-1} X_{-k} (\Sigma_x)_{-k,-k}^{-1} X_{-k}^\top A_0(X, \lambda)^{-1} \right) \\
& \leq C \sigma^2 t \mu_1 \left(A_k(X, \lambda)^{-1} \right)^2 \text{Tr} \left(X_{-k} (\Sigma_x)_{-k,-k}^{-1} X_{-k}^\top \right),
\end{aligned}$$

where we used $\mu_1 \left(A_k(X, \lambda)^{-1} \right) \geq \mu_1 \left(A_0(X, \lambda)^{-1} \right)$. Combining these two inequalities and (E.4) gives the result. □

For $W \in \mathbb{R}^{d \times d}$ which corresponds to eigenvalues of Σ_x , we can obtain Corollary 1 by applying Proposition 2 with $X = XU$.

F EQUIVALENCE OF UPPER BOND AND LOWER BOUND

Proposition 3 (Tsigler and Bartlett (2020)). *Let,*

$$\begin{aligned} \underline{B} &:= \sum_{i=1}^d \frac{\lambda_i \sigma_i^2}{\left(1 + \frac{\lambda_i}{\lambda_{k+1} \rho_k}\right)^2}, & \bar{B} &:= \sum_{i=1}^k \lambda_i \tilde{\sigma}_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \tilde{\sigma}_i^2, \\ \underline{V} &:= \frac{1}{n} \sum_{i=1}^d \min \left\{ 1, \frac{\lambda_i^2}{\sigma_x^4 \lambda_{k+1}^2 (\rho_k + 2)^2} \right\}, & \bar{V} &= \frac{k}{n} + \frac{1}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2}. \end{aligned}$$

Suppose $\rho_k \in (a, b)$ for some $b > a > 0$. Then, it holds that

$$1 \leq \frac{\bar{B}}{\underline{B}} \leq \max \left\{ (1+b)^2, (1+a^{-1})^2 \right\}, \quad 1 \leq \frac{\bar{V}}{\underline{V}} \leq \max \left\{ (2+b)^2, (1+2a^{-1})^2 \right\}.$$

Alternatively, if $k = \min\{l : \rho_l > b\}$ and $b > 1/n$, then, it holds that

$$1 \leq \frac{\bar{B}}{\underline{B}} \leq \max \left\{ (1+b)^2, (1+b^{-1})^2 \right\}, \quad 1 \leq \frac{\bar{V}}{\underline{V}} \leq \max \left\{ (2+b)^2, (1+2b^{-1})^2 \right\}.$$

G EQUIVALENCE OF PROPOSED MODEL AND NORMAL RIDGE REGRESSION

Proposition 4. *Let $\tilde{B}_{\text{Norm}} = \sum_{i=1}^k \lambda_i \tilde{\sigma}_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \tilde{\sigma}_i^2$, $\tilde{V}_{\text{Norm}} = \frac{k\sigma^2}{n} + \frac{\sigma^2}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2}$. When for some constant $\alpha > 0$ we can write $\Sigma_\beta = \alpha I_d$, if there exists $k < n$ such that $\mu_{k+1} \leq \frac{\sigma^2}{n} \leq \lambda_k$, then for some $\lambda \in \mathbb{R}$, it holds that*

$$\tilde{B}_{\text{Norm}} + \tilde{V}_{\text{Norm}} \approx \sum_{i=1}^d \min \left\{ \frac{\sigma^2}{n}, \lambda_i \right\} = U_{\text{NN}}.$$

Proof. When it holds $\tilde{\sigma}_i^2 = \alpha$, if we choose k such that $\lambda_{k+1} \leq \frac{\sigma^2}{n} \leq \lambda_k$ and $\lambda \in \mathbb{R}$ such that $\rho_k = \Theta(1)$, then we can evaluate \tilde{B}_{Norm} and \tilde{V}_{Norm} as

$$\begin{aligned} \tilde{B}_{\text{Norm}} &= \sum_{i=1}^k \lambda_i \tilde{\sigma}_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \tilde{\sigma}_i^2 \\ &\approx \sum_{i=1}^k \frac{\lambda_k}{\lambda_i} \frac{\sigma^2}{n} + \sum_{i=k+1}^d \lambda_i, \\ \tilde{V}_{\text{Norm}} &= \frac{k\sigma^2}{n} + \frac{\sigma^2}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2} \\ &\approx \sum_{i=1}^k \frac{\sigma^2}{n} + \sum_{i=k+1}^d \frac{\lambda_i}{\lambda_k}. \end{aligned}$$

Therefore, since $\lambda_{k+1} \leq \frac{\sigma^2}{n} \leq \lambda_k$ it holds

$$\begin{aligned} \tilde{B}_{\text{Norm}} + \tilde{V}_{\text{Norm}} &\approx \frac{\sigma^2}{n} \sum_{i=1}^k \max \left\{ 1, \frac{\lambda_k}{\lambda_i} \right\} + \sum_{i=k+1}^d \lambda_i \max \left\{ 1, \frac{\lambda_i}{\lambda_k} \right\} \\ &= \frac{\sigma^2}{n} \sum_{i=1}^k 1 + \sum_{i=k+1}^d \lambda_i \\ &= \min \left\{ \frac{\sigma^2}{n}, \lambda_i \right\}. \end{aligned}$$

□

H EXAMPLES

H.1 PROOF OF EXAMPLE 1

Proof. In the vanilla ridge regression, for any $k < n$, the bias term can be lower bounded as follows:

$$\begin{aligned} B_{\text{Norm}} &\gtrsim \sum_{i=1}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2 \\ &> \lambda_1 \sigma_1^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_1^2} \\ &= \rho_k^2 \frac{n}{e \log n}. \end{aligned}$$

Therefore, for $B_{\text{Norm}} = o(1)$ it is necessary that $\rho_k^{-2} = \omega\left(\frac{n}{\log n}\right)$. However, in that case, we can evaluate variance V_{Norm} as

$$\begin{aligned} \frac{V_{\text{Norm}}}{L^2 t} &= \frac{k}{n} \sigma^2 + \frac{\sigma^2}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2} \\ &> \frac{\sigma^2}{n} \sum_{i=n}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2} \\ &\gtrsim \frac{n}{n \log n} \sum_{i=n}^d \frac{1}{i+1-n} \\ &= \frac{1}{\log n} \sum_{i=1}^{d+1-n} \frac{1}{i} \\ &> \frac{1}{\log n} \log(d+1-n) \\ &> \frac{\log(n+1)}{\log n} \\ &> 1. \end{aligned}$$

On the other hand, the obtained two-layer neural network achieves a predictive risk converging to zero:

$$\begin{aligned} U_{\text{NN}} &= \sum_{i=1}^d \min \left\{ \lambda_i \sigma_i^2, \frac{\sigma'^2}{n} \right\} \\ &= \sum_{i=1}^n \min \left\{ \lambda_i \sigma_i^2, \frac{\sigma'^2}{n} \right\} + \sum_{i=n+1}^d \min \left\{ \lambda_i \sigma_i^2, \frac{\sigma'^2}{n} \right\} \\ &= \sum_{i=1}^n \min \left\{ \frac{n}{e^i \log n}, \frac{1}{n} \right\} + \sum_{i=n+1}^d \min \left\{ \frac{n}{(i+1-n)e^i \log n}, \frac{1}{n} \right\} \\ &< \sum_{i=1}^n \min \left\{ \frac{n}{e^i}, \frac{1}{n} \right\} + \sum_{i=n+1}^d \min \left\{ \frac{n}{e^i}, \frac{1}{n} \right\} \\ &= \sum_{i=1}^{\lfloor 2 \log n \rfloor} \frac{1}{n} + \sum_{i=\lfloor 2 \log n \rfloor}^d \frac{n}{e^i} \\ &\leq \frac{2 \log n}{n} + \frac{n}{n^2} \frac{1 - e^{-d-1+\lceil 2 \log n \rceil}}{1 - \frac{1}{e}} \\ &\lesssim \frac{\log n}{n}. \end{aligned}$$

□

H.2 PROOF OF EXAMPLE 2

Proof. In the vanilla ridge regression, for any $k < n$, the bias term can be lower bounded as follows:

$$\begin{aligned} \frac{B_{\text{Norm}}}{L^4} &= \sum_{i=1}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2 \\ &> \lambda_d \sigma_d^2 \\ &= e^{d-d} \\ &= 1. \end{aligned}$$

Therefore, the bias term does not vanish as n goes to ∞ . On the other hand, the obtained two-layer neural network achieves a predictive risk converging to zero:

$$\begin{aligned} U_{\text{NN}} &= \sum_{i=1}^d \min \left\{ \lambda_i \sigma_i^2, \frac{\sigma_i^2}{n} \right\} \\ &= \sum_{i=1}^d \min \left\{ e^{i-d}, \frac{1}{n} \right\} \\ &= \frac{1}{e^d} \sum_{i=1}^{\lfloor d-\log n \rfloor} e^i + \sum_{i=\lceil d-\log n \rceil}^d \frac{1}{n} \\ &\leq \frac{1}{e^{d-1}(e-1)} \left(\frac{e^d}{n} - 1 \right) + \frac{\log n}{n} \\ &\lesssim \frac{1}{n} + \frac{\log n}{n} \\ &\lesssim \frac{\log n}{n}. \end{aligned}$$

□

H.3 PROOF OF EXAMPLE 3

Proof. In the vanilla ridge regression, it holds that

$$\begin{aligned} \frac{V_{\text{Norm}}}{L^2 t} &= \frac{k}{n} \sigma^2 + \frac{\sigma^2}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2} \\ &> \frac{\sigma^2}{n \rho_k^2 \lambda_{k+1}^2} \sum_{i=n+1}^d \lambda_i^2 \\ &= \frac{\sigma^2 (d-n)}{n^3 \rho_k^2 \lambda_{k+1}^2}. \end{aligned}$$

Therefore, $\rho_k^2 \lambda_{k+1}^2 = \omega\left(\frac{\sigma^2 (d-n)}{n^3}\right)$ is necessary for the variance term go to zero. However, in that situation, it holds that

$$\begin{aligned} \frac{B_{\text{Norm}}}{L^4} &= \sum_{i=1}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2 \\ &= \rho_k^2 \lambda_{k+1}^2 \sum_{i=1}^k \frac{\sigma_i^2}{\lambda_i} + \sum_{i=k+1}^d \lambda_i \sigma_i^2 \\ &> \frac{\sigma^2 (d-n)}{n^3} \frac{\sigma_1^2}{\lambda_1} \\ &= 1 - \frac{1}{n^2} \end{aligned}$$

$$\gtrsim 1.$$

Therefore, the bias term doesn't vanish. On the contrary, in obtained two-layer neural network, it holds that

$$\begin{aligned}
U_{\text{NN}} &= \sum_{i=1}^d \min \left\{ \lambda_i \sigma_i^2, \frac{\sigma_i^2}{n} \right\} \\
&= \sum_{i=1}^n \min \left\{ i^{-3}, \frac{1}{n} \right\} + \sum_{i=n+1}^d \min \left\{ n^{-1} i^{-2}, \frac{1}{n} \right\} \\
&= \sum_{i=1}^{\lfloor n^{\frac{1}{3}} \rfloor} \frac{1}{n} + \sum_{i=\lceil n^{\frac{1}{3}} \rceil}^n i^{-3} + \sum_{i=n+1}^d \frac{1}{ni^2} \\
&\leq \frac{n^{\frac{1}{3}}}{n} + \int_{\lfloor n^{\frac{1}{3}} \rfloor}^n x^{-3} dx + \frac{1}{n} \int_n^d x^{-2} dx \\
&= n^{-\frac{2}{3}} + \frac{1}{2} \left(\lfloor n^{\frac{1}{3}} \rfloor \right)^2 - \frac{1}{2} n^{-2} + \frac{1}{n^2} - \frac{1}{nd} \\
&\lesssim \frac{3}{2n^{\frac{2}{3}}} - \frac{1}{nd} \\
&\lesssim n^{-\frac{2}{3}}.
\end{aligned}$$

□

H.4 PROOF OF EXAMPLE 4

Proof. For the vanilla ridge regression, if $k < n^{\frac{2}{3}}$, it holds

$$\begin{aligned}
\frac{B_{\text{Norm}}}{L^4} &= \sum_{i=1}^k \lambda_i \tilde{\sigma}_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \tilde{\sigma}_i^2 \\
&\geq \lambda_{n^{\frac{2}{3}}} \sigma_{n^{\frac{2}{3}}}^2 \\
&= 1.
\end{aligned}$$

On the other hand, if $n^{\frac{2}{3}} \leq k$, it holds

$$\begin{aligned}
\frac{B_{\text{Norm}}}{L^4} &= \sum_{i=1}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2 \\
&= \sum_{i=1}^{n^{\frac{2}{3}}} \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=n^{\frac{2}{3}+1}}^k \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} + \sum_{i=k+1}^d \lambda_i \sigma_i^2 \\
&> \sum_{i=1}^{n^{\frac{2}{3}}} \lambda_i \sigma_i^2 \frac{\rho_k^2 \lambda_{k+1}^2}{\lambda_i^2} \\
&= \rho_k^2 \lambda_{k+1}^2 \sum_{i=1}^{n^{\frac{2}{3}}} \frac{\sigma_i^2}{\lambda_i} \\
&= \rho_k^2 \lambda_{k+1}^2 n^{\frac{3}{2}} \sum_{i=1}^{n^{\frac{2}{3}}} i \\
&\geq \frac{\rho_k^2 \lambda_{k+1}^2}{2} n^{\frac{5}{2}}.
\end{aligned}$$

$\rho_k^2 \lambda_{k+1}^2 = o\left(n^{-\frac{5}{2}}\right)$ is necessary for bias term to vanish. However, in this situation, it holds

$$\begin{aligned} \frac{V_{Norm}}{L^2 t} &= \frac{k}{n} \sigma^2 + \frac{\sigma^2}{n} \sum_{i=k+1}^d \frac{\lambda_i^2}{\rho_k^2 \lambda_{k+1}^2} \\ &> n^{\frac{3}{2}} \sum_{i=n+1}^d i^{-2} \\ &> n^{\frac{3}{2}} \int_n^d x^{-2} dx \\ &= n^{\frac{3}{2}} \left(\frac{1}{n} - \frac{1}{d} \right). \end{aligned}$$

In the obtained two-layer neural network, it holds

$$\begin{aligned} U_{NN} &= \sum_{i=1}^d \min \left\{ \lambda_i \sigma_i^2, \frac{\sigma_i^2}{n} \right\} \\ &= \sum_{i=1}^{n^{\frac{2}{3}}} \min \left\{ n^{\frac{3}{2}} i^{-1}, \frac{1}{n} \right\} + \sum_{i=n^{\frac{2}{3}}}^d \min \left\{ i^{-2}, \frac{1}{n} \right\} \\ &< \frac{n^{\frac{2}{3}}}{n} + \int_{n^{\frac{2}{3}}}^d x^{-2} dx \\ &= n^{-\frac{1}{3}} + n^{-\frac{2}{3}} - \frac{1}{d} \\ &\lesssim \frac{1}{n^{1/3}} - \frac{1}{d} \\ &< n^{-\frac{1}{3}}. \end{aligned}$$

□

I AUXILIARY LEMMAS

Lemma 10. *Suppose $Z \in \mathbb{R}^d$ is a matrix with independent isotropic sub-gaussian rows with $\|Z\|_{\psi_2} \leq \sigma$. Consider $\Sigma \in \mathbb{R}^{d \times d}$ is a positive semidefinite matrix. Then it holds that $\left\| (Z\Sigma)_{k,k} \Sigma_{k,k}^{-1} \right\|_{\psi_2} \leq \sigma$ and $\left\| (Z\Sigma)_{-k,-k} \Sigma_{-k,-k}^{-1} \right\|_{\psi_2} \leq \sigma$.*

Proof.

$$\begin{aligned} \sup_{\|s_k\|=1} \langle s_k, (Z\Sigma)_{k,k} \Sigma_{k,k}^{-1} \rangle &= \sup_{\|s_k\|=1} \left\langle \begin{pmatrix} s_k \\ 0 \end{pmatrix}, Z\Sigma \begin{pmatrix} \Sigma_{k,k}^{-1} & 0 \\ 0 & 0 \end{pmatrix} \right\rangle \\ &= \sup_{\|s_k\|=1} \left\langle \begin{pmatrix} s_k \\ 0 \end{pmatrix}, Z \begin{pmatrix} I_k & 0 \\ 0 & 0 \end{pmatrix} \right\rangle \\ &= \sup_{\|s_k\|=1} \left\langle \begin{pmatrix} s_k \\ 0 \end{pmatrix}, \begin{pmatrix} Z_k \\ 0 \end{pmatrix} \right\rangle \\ &= \sup_{\|s_k\|=1} \left\langle \begin{pmatrix} s_k \\ 0 \end{pmatrix}, Z \right\rangle \\ &\leq \sup_{\|s\|=1} \langle s, Z \rangle. \end{aligned}$$

Therefore, it holds $\left\| (Z\Sigma)_{k,k} \Sigma_{k,k}^{-1} \right\|_{\psi_2} \leq \sigma$. We can show the second assertion, in the same way as the first assertion. □

Lemma 11. Suppose $Z \in \mathbb{R}^{n \times d}$ is a matrix with independent isotropic sub-gaussian rows with $\|Z\|_{\psi_2} \leq \sigma$. Consider $\Sigma \in \mathbb{R}^{d \times d}$ is a positive semidefinite matrix. Then for some absolute constant c for any $t \in (0, n)$ and $0 < k < d$, it holds that with probability at least $1 - 2e^{-ct}$,

$$(n - t\sigma^2) \text{Tr}(\Sigma_{-k, -k}) \leq \sum_{i=1}^n \left\| \left(\Sigma^{1/2} Z_i^\top \right)_{-k} \right\|^2 \leq (n + t\sigma^2) \text{Tr}(\Sigma_{-k, -k}).$$

Proof. Let eigenvectors of $\Sigma_{-k, -k}$ as $U_{-k} \in \mathbb{R}^{d \times d}$. Then,

$$\left| \sum_{i=1}^n \left\| \left(\Sigma^{1/2} Z_i^\top \right)_{-k} \right\|^2 - \text{Tr}(\Sigma_{-k, -k}) \right| = \left| \sum_{i=1}^n \left\| U_{-k}^\top \left(\Sigma^{1/2} Z_i^\top \right)_{-k} \right\|^2 - \sum_{i=1}^d \mu_i(\Sigma) \right|.$$

From Lemma 10, it holds that $\left\| (Z\Sigma)_{-k, -k} \Sigma_{-k, -k}^{-1} \right\|_{\psi_2} \leq \sigma$. Therefore, applying Lemma 17 in (Tsigler and Bartlett, 2020) gives the result. \square

Lemma 12. Let x_1, \dots, x_n be i.i.d. centered random vectors in \mathbb{R}^l with covariance matrix Σ . If x is sub-Gaussian with $\|x\|_{\psi_2} \leq \sigma_x$, then there exist constants C and c which only depends on σ_x such that it holds that with probability at least $1 - 2\exp(-ct^2)$,

$$\left(1 - C\sqrt{\frac{l}{n}} - t \right) \Sigma \leq \hat{\Sigma} \leq \left(1 + C\sqrt{\frac{l}{n}} + t \right) \Sigma.$$

Proof. By Theorem 5.39 in (Vershynin, 2012), there exist C' and c' which only depends on σ_x such that it holds that with probability at least $1 - 2e^{-c't^2}$,

$$\begin{aligned} \|I_l - \Sigma^{-\frac{1}{2}} \hat{\Sigma} \Sigma^{-\frac{1}{2}}\|_{\text{op}} &\leq 2C\sqrt{\frac{l}{n}} + 2t, \\ \therefore -\left(2C'\sqrt{\frac{l}{n}} + 2t \right) I_l &\leq I_l - \Sigma^{-\frac{1}{2}} \hat{\Sigma} \Sigma^{-\frac{1}{2}} \leq \left(2C'\sqrt{\frac{l}{n}} + 2t \right) I_l, \\ \therefore \left(1 - 2C'\sqrt{\frac{l}{n}} - 2t \right) \Sigma &\leq \hat{\Sigma} \leq \left(2C'\sqrt{\frac{l}{n}} - 2t \right) \Sigma. \end{aligned}$$

\square

Lemma 13. Suppose $A \in \mathbb{R}^{n \times n}$ is a positive semidefinite matrix and $\epsilon^{(i)}$ ($i = 1, \dots, m$) are centered vectors whose components are independent and have sub-gaussian norm at most σ . Then for some constants c, C , for any $t > 1$ it holds that with probability at least $1 - 2e^{-t/c}$,

$$\frac{1}{m} \sum_{i=1}^m \epsilon^{(i)\top} A \epsilon^{(i)} \leq C\sigma^2 t \text{Tr}(A).$$

Proof. From Lemma 18 in Tsigler and Bartlett (2020) for any i there exists constants C_ϵ, c_i it holds that with probability at least $1 - 2e^{-t/c_i}$,

$$\epsilon^{(i)\top} A \epsilon^{(i)} \leq C_\epsilon \sigma^2 t \text{Tr}(A).$$

Notice that C_ϵ doesn't depend on the choice of i . Let $c := \min_i c_i$ and $J := \arg\max_i \epsilon^{(i)\top} A \epsilon^{(i)}$, and we have

$$\begin{aligned} \Pr\left(\epsilon^{(J)\top} A \epsilon^{(J)} \leq C_\epsilon \sigma^2 t \text{Tr}(A) \right) &> 1 - 2e^{-t/c_J} \\ &> 1 - 2e^{-t/c}. \end{aligned}$$

Therefore it holds that with probability at least $1 - 2e^{-t/c}$,

$$\frac{1}{m} \sum_{i=1}^m \epsilon^{(i)\top} A \epsilon^{(i)} \leq \left(\epsilon^{(J)} \right)^\top A \epsilon^{(J)} \leq C_\epsilon \sigma^2 t \text{Tr}(A).$$

\square