Risk Bounds for Over-parameterized Maximum Margin Classification on Sub-Gaussian Mixtures

Yuan Cao* and Quanquan Gu† and Mikhail Belkin‡

Abstract

Modern machine learning systems such as deep neural networks are often highly over-parameterized so that they can fit the noisy training data exactly, yet they can still achieve small test errors in practice. In this paper, we study this "benign overfitting" (Bartlett et al. (2020)) phenomenon of the maximum margin classifier for linear classification problems. Specifically, we consider data generated from sub-Gaussian mixtures, and provide a tight risk bound for the maximum margin linear classifier in the over-parameterized setting. Our results precisely characterize the condition under which benign overfitting can occur in linear classification problems, and improve on previous work. They also have direct implications for over-parameterized logistic regression.

1 Introduction

In modern machine learning, complex models such as deep neural networks have received increasing popularity. These complicated models are known to be able to fit noisy training data sets, while at the same time achieving small test errors. In fact, this *benign overfitting* phenomenon is not a unique feature of deep learning. Even for kernel methods and linear models, Belkin et al. (2018) demonstrated that interpolators on the noisy training data can still perform near optimally on the test data. A series of recent works (Belkin et al., 2019b; Muthukumar et al., 2020b; Hastie et al., 2019; Bartlett et al., 2020) theoretically studied how over-parameterization can achieve small population risk.

In particular in Bartlett et al. (2020) the authors considered the setting where the data are generated from a ground-truth linear model with noise, and established a tight population risk bound for the minimum norm linear interpolator with a matching lower bound. More recently, Tsigler and Bartlett (2020) further studied benign overfitting in ridge regression, and established non-asymptotic generalization bounds for over-parametrized ridge regression. They showed that those bounds are tight for a range of regularization parameter values. Notably, these results cover arbitrary covariance structure of the data, and give a nice characterization of how the spectrum of the data covariance matrix affects the population risk in the over-parameterized regime.

*Department of Computer Science, University of California, Los Angeles, CA 90095, USA; e-mail: yuancao@cs.ucla.edu

†Department of Computer Science, University of California, Los Angeles, CA 90095, USA; e-mail: qgu@cs.ucla.edu

‡Halicioğlu Data Science Institute, University of California, San Diego, La Jolla, CA 92093, USA; e-mail: mbelkin@ucsd.edu

Very recently, benign overfitting has also been studied in the setting of linear classification (Chatterji and Long, 2020; Muthukumar et al., 2020a; Wang and Thrampoulidis, 2020). Specifically, Muthukumar et al. (2020a) studied the setting where the data inputs are Gaussian and the labels are generated from a ground truth linear model with label flipping noise, and showed equivalence between the hard-margin support vector machine (SVM) solution and the minimum norm interpolator to study benign overfitting. Chatterji and Long (2020); Wang and Thrampoulidis (2020) studied the benign overfitting phenomenon in sub-Gaussian/Gaussian mixture models and established population risk bounds for the maximum margin classifier. Chatterji and Long (2020) leveraged the *implicit bias* of gradient descent for logistic regression (Soudry et al., 2018) to establish the risk bound. Wang and Thrampoulidis (2020) established an equivalence result between classification and regression for isotropic Gaussian mixture models. While these results have offered valuable insights into the benign overfitting phenomenon for (sub-)Gaussian mixture classification, they still have certain limitations. Unlike the results in the regression setting where the eigenvalues of the data covariance matrix play a key role, the current results for Gaussian/sub-Gaussian mixture models do not show the impact of the spectrum of the data covariance matrix on the risk.

In this paper, we study the benign overfitting phenomenon in a general sub-Gaussian mixture model that covers both the isotropic and anisotropic settings, where the d-dimensional features from two classes have the same covariance matrix $\mathbf{\Sigma}$ but have different means $\boldsymbol{\mu}$ and $-\boldsymbol{\mu}$ respectively. We consider the over-parameterized setting where d is larger than the sample size n, and prove a risk bound for the maximum margin classifier. Our proof of the risk bound has a key step to demonstrate that under certain conditions regarding the eigenvalues of $\mathbf{\Sigma}$, the mean vector $\boldsymbol{\mu}$ and the sample size n, the maximum margin classifier for this problem is identical to the minimum norm interpolator. We then utilize this result to establish a tight population risk bound of the maximum margin classifier. Our result reveals how the eigenvalues of the covariance matrix $\mathbf{\Sigma}$ affect the benign property of the classification problem, and is tighter and more general than existing results on sub-Gaussian/Gaussian mixture models. The contributions of this paper are as follows:

- We establish a tight population risk bound for the maximum margin classifier. Our bound works for both the isotropic and anisotropic settings, which is more general than existing results in Chatterji and Long (2020); Wang and Thrampoulidis (2020). When reducing our bound to the setting studied in Chatterji and Long (2020), our result gives a bound $\exp(-\Omega(n\|\boldsymbol{\mu}\|_2^4/d))$, where n is the training sample size. Our bound is tighter than the risk bound $\exp(-\Omega(\|\boldsymbol{\mu}\|_2^4/d))$ in Chatterji and Long (2020) by a factor of n in the exponent. Our result also gives a tighter risk bound than that in Wang and Thrampoulidis (2020) in the so-called "low SNR setting": our result suggests that $\|\boldsymbol{\mu}\|_2^4 = \omega(d/n)$ suffices to ensure an $o(1)$ population risk, while Wang and Thrampoulidis (2020) requires $\|\boldsymbol{\mu}\|_2^4 = \omega((d/n)^{3/2})$.

- We establish population risk lower bounds achieved by the maximum margin classifier under two different settings. In both settings, the lower bounds match our population risk upper bound up to some absolute constants. This suggests that our population risk bound is tight.

- Our analysis reveals that for a class of high-dimensional anisotropic sub-Gaussian mixture models, the maximum margin linear classifier on the training data can achieve small population risk under mild assumptions on the sample size n and mean vector $\boldsymbol{\mu}$. Specifically, suppose that the eigenvalues of $\mathbf{\Sigma}$ are $\{\lambda_k = k^{-\alpha}\}_{k=1}^d$ for some parameter $\alpha \in [0, 1)$, and treat the sample size n as a constant. Then our result shows that to achieve $o(1)$ population risk, the following conditions on $\|\boldsymbol{\mu}\|_2$ suffice:

$$\|\boldsymbol{\mu}\|_2 = \begin{cases} \omega(d^{1/4-\alpha/2}), & \text{if } \alpha \in [0, 1/2), \\ \omega((\log(d))^{1/4}), & \text{if } \alpha = 1/2, \\ \omega(1). & \text{if } \alpha \in (1/2, 1). \end{cases}$$

More specifically, when $\alpha = 1/2$, the condition on the mean vector $\boldsymbol{\mu}$ only has a logarithmic dependency on the dimension d, and when $\alpha \in (1/2, 1)$, the condition on $\boldsymbol{\mu}$ for benign overfitting is dimension free.

- Our proof of the population risk bound introduces some tight intermediate results, which may be of independent interest. Specifically, our proof utilizes polarization identity to establish equivalence between the maximum margin classifier and the minimum norm interpolator. This is, to the best of our knowledge, the first equivalence result between classification and regression for anisotropic sub-Gaussian mixture models.

1.1 Additional Related Work

Benign overfitting is closely related to the phenomenon of double descent studied in recent works. Belkin et al. (2019a,b) showed experimental results and provided theoretical analyses on some specific models to demonstrate that the risk curve versus over-parameterization has a double descent shape. These results can therefore indicate that over-parameterization can be beneficial to achieve small test risk. Hastie et al. (2019); Wu and Xu (2020) studied the double descent phenomenon in linear regression under the setting where the dimension d and sample size n can grow simultaneously but have a fixed ratio, and showed that the population risk exhibits a double descent curve with respect to the ratio. More recently, Mei and Montanari (2019); Liao et al. (2020); Montanari and Zhong (2020) further extended the setting to random feature models and studied double descent when the sample size, data dimension and the number of random features have fixed ratios.

Our work is also related to the studies of implicit bias, which analyze the impact of training algorithms when the over-parameterized models have multiple global minima. Specifically, Soudry et al. (2018) showed that if the training data are linearly separable, then gradient descent on unregularized logistic regression converges directionally to the maximum margin linear classifier on the training data set. Ji and Telgarsky (2019) further studied the implicit bias of gradient descent for logistic regression on non-separable data. Gunasekar et al. (2018) studied the implicit bias of various optimization methods for generic objective functions. Gunasekar et al. (2017); Arora et al. (2019) established implicit bias results for matrix factorization problems. More recently, Lyu and Li (2019) showed that gradient flow for learning homogeneous neural networks with logistic loss maximizes the normalized margin on the training data set. These studies of implicit bias offer a handle for us to connect the over-parameterized logistic regression with the maximum margin classifiers for linear models.

2 Problem Setting

In this section, we introduce the notations and detailed problem setup. We begin with the notations.

2.1 Notations

We use lower case letters to denote scalars, and use lower/upper case bold face letters to denote vectors/matrices respectively. For a vector $\mathbf{v}$, we denote by $\|\mathbf{v}\|_2$ the ℓ_2-norm of $\mathbf{v}$. For a matrix $\mathbf{A}$, we use $\|\mathbf{A}\|_2$, $\|\mathbf{A}\|_F$ to denote its spectral norm and Frobinuous norm respectively, and use $\mathrm{tr}(\mathbf{A})$ to denote its trace. For a vector $\mathbf{v} \in \mathbb{R}^d$ and a positive definite matrix $\mathbf{A}$, we define $\|\mathbf{v}\|_{\mathbf{A}} = \sqrt{\mathbf{v}^\top \mathbf{A} \mathbf{v}}$. For an integer n, we denote $[n] = \{1, 2, \dots, n\}$.

We also use standard asymptotic notations $O(\cdot)$, $\Omega(\cdot)$, $o(\cdot)$, and $\omega(\cdot)$. Let $\{a_n\}$ and $\{b_n\}$ be two sequences. If there exists a constant $C > 0$ such that $|a_n| \leq C|b_n|$ for all large enough n, then we denote $a_n = O(b_n)$. We denote $a_n = \Omega(b_n)$ if $b_n = O(a_n)$. Moreover, we write $a_n = o(b_n)$ if $\lim |a_n/b_n| = 0$ and $a_n = \omega(b_n)$ if $\lim |a_n/b_n| = \infty$. We also use $\widetilde{O}(\cdot)$ and $\widetilde{\Omega}(\cdot)$ to hide some logarithmic terms in Big-O and Big-Omega notations.

At last, for a random variable Z, we denote by $\|Z\|_{\psi_2}$ and $\|Z\|_{\psi_1}$ the sub-Gaussian and sub-exponential norms of Z respectively.

2.2 Sub-Gaussian Mixture Model

We consider a model where the feature vectors are generated from a mixture of two sub-Gaussian distributions with means $\boldsymbol{\mu}$ and $-\boldsymbol{\mu}$ and the same covariance matrix $\boldsymbol{\Sigma}$. We assume that each data pair $(\mathbf{x}, y)$ are generated independently from the following procedure:

1. The label $y \in \{+1, -1\}$ is generated as a Rademacher random variable.

2. A random vector $\mathbf{u} \in \mathbb{R}^d$ is generated from a distribution such that the entries of $\mathbf{u}$ are independent sub-Gaussian random variables with $\mathbb{E}[u_j] = 0$, $\mathbb{E}[u_j^2] = 1$ and $\|u_j\|_{\psi_2} \leq \sigma_u$ for all $j \in [d]$.

3. Let $\boldsymbol{\Sigma}$ be a positive definite matrix with eigenvalue decomposition $\boldsymbol{\Sigma} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top$, where $\boldsymbol{\Lambda} = \mathrm{diag}\{\lambda_1, \dots, \lambda_d\}$ and $\mathbf{V}$ is an orthonormal matrix consisting of the eigenvectors of $\boldsymbol{\Sigma}$. We calculate the random vector $\mathbf{q}$ based on $\mathbf{u}$ as $\mathbf{q} = \mathbf{V}\boldsymbol{\Lambda}^{1/2}\mathbf{u}$. This ensures that $\mathbf{q}$ has mean zero and a covariance matrix $\boldsymbol{\Sigma}$.

4. The feature is given as $\mathbf{x} = y \cdot \boldsymbol{\mu} + \mathbf{q}$, where $\boldsymbol{\mu} \in \mathbb{R}^d$ is a vector. Clearly, the mean of $\mathbf{x}$ is $\boldsymbol{\mu}$ when $y = 1$ and is $-\boldsymbol{\mu}$ when $y = -1$.

We consider n training data points $(\mathbf{x}_i, y_i)$ generated independently from the above procedure, and denote

$$\mathbf{X} = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q},$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top, \mathbf{Q} = [\mathbf{q}_1, \dots, \mathbf{q}_n]^\top \in \mathbb{R}^{n \times d}$, and $\mathbf{y} = [y_1, \dots, y_n]^\top \in \{\pm 1\}^n$. For any $\boldsymbol{\theta} \in \mathbb{R}^d$, the population risk of the linear classifier $\mathbf{x} \to \langle \boldsymbol{\theta}, \mathbf{x} \rangle$ is defined as:

$$R(\boldsymbol{\theta}) = \mathbb{P}\big(y \cdot \langle \boldsymbol{\theta}, \mathbf{x} \rangle < 0\big).$$

In this paper, we consider the maximum margin linear classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$, i.e., the solution to the hard-margin support vector machine:

$$\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \mathrm{argmin}\, \|\boldsymbol{\theta}\|_2^2, \quad \text{subject to } y_i \cdot \langle \boldsymbol{\theta}, \mathbf{x}_i \rangle \geq 1, i \in [n],$$

and study its population risk $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$.

A recent work (Chatterji and Long, 2020) has studied a similar sub-Gaussian mixture model under an assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) = \Omega(d)$, and considered additional label flipping noises. In this paper, we do not introduce the label flipping noises for simplicity, but we consider a general covariance matrix $\boldsymbol{\Sigma}$ to cover the general anisotropic setting. It is worth noting that although our model is not exactly the same as Chatterji and Long (2020) because we don't have additional label flipping noise, there is still noise in our model because of the nature of sub-Gaussian mixture model. For example, consider a mixture of two Gaussian distributions. The two Gaussian clusters have non-trivial overlap, and the Bayes optimal classifier has non-zero Bayes risk. Therefore, the Bayes optimal classifier and the interpolating classifier are generally quite different. In general, a model is appropriate for the study of benign overfitting whenever the optimal classifier has non-zero Bayes risk.

Our model is rather general and covers the following examples.

Example 2.1 (Gaussian mixture model)**.** The most straight-forward example is when the data are generated from Gaussian mixtures $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $N(-\boldsymbol{\mu}, \boldsymbol{\Sigma})$. This is covered by our model when the sub-Gaussian vector $\mathbf{u}$ is a standard Gaussian random vector.

Example 2.2 (Rare/weak feature model)**.** The rare-weak model is a special case of the Gaussian mixture model where $\boldsymbol{\Sigma} = \mathbf{I}$ and $\boldsymbol{\mu}$ is a sparse vector with s non-zero entries equaling γ.

The rare/weak feature model was originally investigated by Donoho and Jin (2008); Jin (2009), and has been recently studied by Chatterji and Long (2020).

2.3 Connection to Over-parameterized Logistic Regression

Our study of the maximum margin classifier is closely related to over-parameterized logistic regression. In logistic regression, we consider the following empirical loss minimization problem:

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} L(\boldsymbol{\theta}) := \frac{1}{n} \sum_{i=1}^{n} \log[1 + \exp(-y_i \cdot \langle \boldsymbol{\theta}, \mathbf{x}_i \rangle)].$$

We solve the above optimization problem with gradient descent

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \cdot \nabla L(\boldsymbol{\theta}^{(t)}), \tag{2.1}$$

where $\eta > 0$ is the learning rate.

In the over-parameterized setting where $d \gg n$, it is evident that the training data points are linearly separable with high probability (for example, $\mathbf{X}\mathbf{X}^\top$ is invertible with high probability and the minimum norm interpolator $\widehat{\boldsymbol{\theta}}_{\mathrm{LS}} = \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y}$ separates the training data.). For linearly separable data, a series of recent works have studied the *implicit bias* of (stochastic) gradient descent for logistic regression (Soudry et al., 2018; Ji and Telgarsky, 2019; Nacson et al., 2019). These results demonstrate that among all linear classifiers that can classify the training data correctly, gradient descent will converge to the one that maximizes the ℓ_2 margin. Such an implicit bias result is summarized in the following lemma.

Lemma 2.3 (Theorem 3 in Soudry et al. (2018))**.** Suppose that the training data set $\{(\mathbf{x}_i, y_i)\}$ is linearly separable. Then as long as $\eta > 0$ is small enough, the gradient descent iterates $\boldsymbol{\theta}^{(t)}$ for logistic regression defined in (2.1) has the following direction limit:

$$\lim_{t\to\infty} \frac{\boldsymbol{\theta}^{(t)}}{\|\boldsymbol{\theta}^{(t)}\|_2} = \frac{\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}}{\|\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}\|_2},$$

where $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ is the maximum margin classifier.

Lemma 2.3 suggests that our risk bound of the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ directly implies a risk bound for the over-parameterized logistic regression trained by gradient descent.

3 Main Results

In this section, we present our main result on the population risk bound of the maximum margin classifier, and then give a lower bound result to demonstrate the tightness of our upper bounds. We also showcase the application of our results to isotropic and anisotropic sub-Gaussian mixture models to study the conditions under which benign overfitting occurs.

The main result of this paper is given in the following theorem, where we establish the population risk bound for the maximum margin classifier $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$.

Theorem 3.1. Suppose that

$$\mathrm{tr}(\boldsymbol{\Sigma}) \geq C \max\left\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\sqrt{\log(n)}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\right\}$$

and $\|\boldsymbol{\mu}\|_2^2 \geq C\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ for some absolute constant C. Then with probability at least $1-n^{-1}$, the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ has the following risk bound

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\left(\frac{-C'n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2}\right),$$

where C' is an absolute constant.

Theorem 3.1 gives the population risk bound of the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$. Based on the implicit bias of gradient descent for over-parameterized logistic regression (Lemma 2.3), we have that the gradient descent iterates $\boldsymbol{\theta}^{(t)}$ satisfy that

$$\lim_{t\to\infty} R(\boldsymbol{\theta}^{(t)}) = \lim_{t\to\infty} R(\boldsymbol{\theta}^{(t)}/\|\boldsymbol{\theta}^{(t)}\|_2) = R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}/\|\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}\|_2) = R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}).$$

Therefore, the same risk bound in Theorem 3.1 also applies to the over-parameterized logistic regression trained by gradient descent.

3.1 Population Risk Lower Bound

In this section we further present lower bounds on the population risk achieved by the maximum margin classifier, which demonstrate that our population risk upper bound in Theorem 3.1 is tight. We have the following theorem.

Theorem 3.2. Consider Gaussian mixture model with covariance matrix $\boldsymbol{\Sigma}$ and mean vectors $\boldsymbol{\mu}$ and $-\boldsymbol{\mu}$. Suppose that $\mathrm{tr}(\boldsymbol{\Sigma}) \geq C \max\left\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\sqrt{\log(n)}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\right\}$, and $\|\boldsymbol{\mu}\|_2^2 \geq C\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$

for some constant C. Then there exist absolute constants C', C'', such that the following results hold:

1. If $n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \geq C(\|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2)$, then with probability at least $1 - n^{-1}$,

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \geq C'' \exp\big(-C'\|\boldsymbol{\mu}\|_2^4/\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2\big).$$

2. If $\|\boldsymbol{\Sigma}\|_F^2 \geq Cn(\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\boldsymbol{\Sigma}\|_2^2)$, then with probability at least $1 - n^{-1}$,

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \geq C'' \exp\big(-C'n\|\boldsymbol{\mu}\|_2^4/\|\boldsymbol{\Sigma}\|_F^2\big).$$

Theorem 3.2 gives lower bounds for the population risk in two settings: (i) $n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \geq C(\|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2)$; and (ii) $\|\boldsymbol{\Sigma}\|_F^2 \geq Cn(\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\boldsymbol{\Sigma}\|_2^2)$. Note that in the population risk upper bound in Theorem 3.1, there are three terms in the denominator of the exponent: $\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2$, $\|\boldsymbol{\Sigma}\|_F^2$, and $n\|\boldsymbol{\Sigma}\|_2^2$. Therefore, setting (i) and setting (ii) in Theorem 3.2 correspond to the cases when the first or the second term is the leading term, respectively. Moreover, it is also easy to check that under both settings, our lower bound in Theorem 3.2 matches the upper bound in Theorem 3.1. This suggests that our population risk bound in Theorem 3.1 is tight.

3.2 Implications for Specific Examples

Theorem 3.1 holds for general covariance matrices $\boldsymbol{\Sigma}$, and illustrates how the spectrum of $\boldsymbol{\Sigma}$ affects the population risk of the maximum margin classifier. This makes our result more general than the recent results in Chatterji and Long (2020); Wang and Thrampoulidis (2020), where the population risk bounds are given only in terms of the sample size n, dimension d and the norm of the mean vector $\|\boldsymbol{\mu}\|_2$. In fact, when we specialize our general result to the isotropic setting, our result also provides a tighter risk bound than these existing results. Specifically, our population risk bound for the isotropic setting is given in the following corollary.

Corollary 3.3 (Isotropic sub-Gaussian mixtures)**.** Consider the setting where $\boldsymbol{\Sigma} = \mathbf{I}$. Suppose that

$$d \geq C \max\big\{n^2, n\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_2\big\}$$

and $\|\boldsymbol{\mu}\|_2 \geq C$ for some absolute constant C. Then with probability at least $1 - n^{-1}$, the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ has the following risk bound

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\bigg(-\frac{C'n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_2^2 + d}\bigg),$$

where C' is an absolute constant.

Remark 3.4. Chatterji and Long (2020) recently gave a risk bound of order $\exp(-\Omega(\|\boldsymbol{\mu}\|_2^4/d))$ for sub-Gaussian mixture models under the condition that $d = \Omega(\max\{n^2\log(n), n\|\boldsymbol{\mu}\|_2^2\})$. Clearly, our result in Corollary 3.3 only requires the condition $d = \Omega(n\|\boldsymbol{\mu}\|_2)$ rather than $d = \Omega(n\|\boldsymbol{\mu}\|_2^2)$. Therefore, our condition on d is milder. Moreover, when the stronger condition $d = \Omega(n\|\boldsymbol{\mu}\|_2^2)$ holds, our risk bound becomes $\exp(-\Omega(n\|\boldsymbol{\mu}\|_2^4/d))$, which is better than the result of Chatterji and Long (2020) by a factor of n in the exponent.

Remark 3.5. Wang and Thrampoulidis (2020) studied the risk bound in two settings, namely the low-SNR setting ($\|\boldsymbol{\mu}\|_2^2 \leq d/n$) and the high-SNR setting ($\|\boldsymbol{\mu}\|_2^2 > d/n$). The numerator $n\|\boldsymbol{\mu}\|_2^2 + d$ in the exponent in Corollary 3.3 provides further insights into the difference between the low- and high-SNR settings. Moreover, for the low-SNR setting, Corollary 2.2 in Wang and Thrampoulidis (2020) requires $\|\boldsymbol{\mu}\|_2^4 = \omega((d/n)^{3/2})$ to make the risk converge to zero. In comparison, Corollary 3.3 only requires $\|\boldsymbol{\mu}\|_2^4 = \omega(d/n)$ to ensure $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) = o(1)$. Therefore, our risk bound is tighter than theirs.

Besides being tighter than previous results when reduced to the isotropic setting, Theorem 3.1 covers both the isotropic and anisotropic settings. In the following, we provide some case studies under the anisotropic setting and show how the decay rate of the eigenvalues of the covariance matrix $\boldsymbol{\Sigma}$ affects the population risk.

It is worth noting that the assumption of Theorem 3.1 requires that $\mathrm{tr}(\boldsymbol{\Sigma})$ is large enough, while the risk bound in Theorem 3.1 only depends on $\|\boldsymbol{\Sigma}\|_F$ and $\|\boldsymbol{\Sigma}\|_2$. In the over-parameterized setting where the dimension d is large, it is possible that for certain covariance matrices $\boldsymbol{\Sigma}$ with appropriate eigenvalue decay rates, $\mathrm{tr}(\boldsymbol{\Sigma}) \gg 1$ while $\|\boldsymbol{\Sigma}\|_F, \|\boldsymbol{\Sigma}\|_2 = O(1)$. This implies that for many anisotropic sub-Gaussian mixture models, the assumptions in Theorem 3.1 can be easily satisfied, while the risk bound can be small at the same time. Following this intuition, we study the conditions under which the the maximum margin interpolator $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ achieves $o(1)$ population risk. We denote by λ_k the k-th largest eigenvalue of $\boldsymbol{\Sigma}$, and consider a polynomial decay spectrum $\{\lambda_k = k^{-\alpha}\}_{k=1}^d$, where we introduce a parameter α to control the eigenvalue decay rate. We have the following corollary.

Corollary 3.6 (Anisotripic sub-Gaussian mixtures with polynomial spectrum decay)**.** Suppose that $\lambda_k = k^{-\alpha}$, n is a large enough constant, and one of the following conditions hold:

1. $\alpha \in [0, 1/2)$, $d = \widetilde{\Omega}((\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$, and $\|\boldsymbol{\mu}\|_2 = \omega(1 + d^{1/4-\alpha/2})$.
2. $\alpha = 1/2$, $d = \widetilde{\Omega}(\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2)$, and $\|\boldsymbol{\mu}\|_2 = \omega((\log(d))^{1/4})$.
3. $\alpha \in (1/2, 1)$, $d = \widetilde{\Omega}((\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$, and $\|\boldsymbol{\mu}\|_2 = \omega(1)$.

Then with probability at least $1 - n^{-1}$, the population risk of the maximum margin classifier satisfies $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) = o(1)$.

Corollary 3.6 follows by calculating the orders of $\mathrm{tr}(\boldsymbol{\Sigma}) = \sum_{k=1}^d \lambda_k$ and $\|\boldsymbol{\Sigma}\|_F^2 = \sum_{k=1}^d \lambda_k^2$. Here we treat the sample size n as a constant for simplicity. A full version of the corollary with detailed dependency on n is given as Corollary C.1 in Appendix C together with the proof. Intuitively, when $\|\boldsymbol{\mu}\|_2$ is large, the two classes are far away from each other and therefore linear classifiers can achieve small population risk. From Corollary 3.6, we can see that the decay rate of the eigenvalues of the covariance matrix $\boldsymbol{\Sigma}$ determines how large $\|\boldsymbol{\mu}\|_2$ needs to be to ensure small population risk: when the $\{\lambda_k\}$ decays faster (i.e., when α is larger), the maximum margin classifier can achieve $o(1)$ population risk with a smaller $\|\boldsymbol{\mu}\|_2$.

Corollary 3.6 also exhibits a certain "phase transition" regarding the eigenvalue decay rate and the conditions on $\|\boldsymbol{\mu}\|_2$. We can see that the eigenvalue decay rate can be divided into three regimes $\alpha \in [0, 1/2)$, $\alpha = 1/2$ and $\alpha \in (1/2, 1)$. Under the condition that $d = \widetilde{\Omega}((\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$, achieving $o(1)$ risk in each of these regimes requires $\|\boldsymbol{\mu}\|_2 = \omega(d^{1/4})$, $\|\boldsymbol{\mu}\|_2 = \omega([\log(d)]^{1/4})$, and $\|\boldsymbol{\mu}\|_2 = \omega(1)$

respectively. Specifically, when $\alpha \in (1/2, 1)$, the condition on $\boldsymbol{\mu}$ is independent of the dimension d. This means that when $\alpha \in (1/2, 1)$, for any $\epsilon > 0$, as long as $\|\boldsymbol{\mu}\|_2 = \Omega(\sqrt{\log(\epsilon)})$, we have

$$\lim_{d\to\infty} R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \epsilon.$$

Therefore, our result covers the infinite dimensional setting when the eigenvalues of the covariance matrix $\boldsymbol{\Sigma}$ have an appropriate decay rate, i.e., $\alpha \in (1/2, 1)$.

Another interesting observation in Theorem 3.1 is that it uses both $\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ and $\|\boldsymbol{\mu}\|_2$, and therefore the alignment between $\boldsymbol{\mu}$ and the eigenvectors of $\boldsymbol{\Sigma}$ can affect the population risk bound. In our discussion above, we have mainly focused on the worst case scenario where the direction of $\boldsymbol{\mu}$ aligns with the first eigenvector of $\boldsymbol{\Sigma}$. In the following corollary, we discuss the case where $\boldsymbol{\mu}$ is parallel to the eigenvector of $\boldsymbol{\Sigma}$ corresponding to the eigenvalue λ_k.

Corollary 3.7 (Risk bounds for $\boldsymbol{\mu}$ along different directions)**.** Suppose that $\boldsymbol{\Sigma}\boldsymbol{\mu} = \lambda_k\boldsymbol{\mu}$ for some $k \in [d]$,

$$\mathrm{tr}(\boldsymbol{\Sigma}) \geq C \max\left\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\sqrt{\lambda_k \log(n)} \cdot \|\boldsymbol{\mu}\|_2\right\}$$

and $\|\boldsymbol{\mu}\|_2^2 \geq C\lambda_k$ for some absolute constant C. Then with probability at least $1 - n^{-1}$, the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ has the following risk bound

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\left(\frac{-C'n\|\boldsymbol{\mu}\|_2^4}{n\lambda_k \cdot \|\boldsymbol{\mu}\|_2^2 + \|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2}\right),$$

where C' is an absolute constant.

We can see that, when $\boldsymbol{\mu}$ aligns with the eigendirections corresponding to a smaller eigenvalue of $\boldsymbol{\Sigma}$, then Corollary 3.7 holds under milder conditions on $\mathrm{tr}(\boldsymbol{\Sigma})$ and $\|\boldsymbol{\mu}\|_2$, and the population risk achieved by the maximum margin solution is also better. This phenomenon perfectly matches the geometric intuition of sub-Gaussian mixture classifications, as is illustrated in Figure 1.

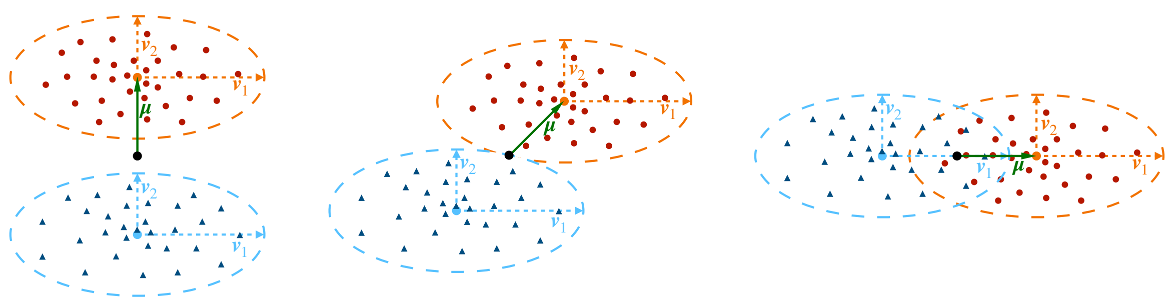

(a) $\boldsymbol{\mu}$ aligns with $\mathbf{v}_2$ (b) $\boldsymbol{\mu}$ points at a random direction (c) $\boldsymbol{\mu}$ aligns with $\mathbf{v}_1$

Figure 1: A 2-dimensional illustration of sub-Gausisan mixture classification problems with different directions of $\boldsymbol{\mu}$. We consider the setting where $\boldsymbol{\Sigma} \in \mathbb{R}^{2\times 2}$ has two eigenvalues $\lambda_1 > \lambda_2$ with the corresponding eigenvectors $\mathbf{v}_1, \mathbf{v}_2$. (a) shows the setting where $\boldsymbol{\mu}$ aligns with $\mathbf{v}_2$. (b) shows the setting where $\boldsymbol{\mu}$ points at a random direction. (c) is for the case when $\boldsymbol{\mu}$ aligns with $\mathbf{v}_1$. These figures clearly show that (a) is the easiest case for classification and (c) is the hardest case. This matches the result in Corollary 3.7.

At last, we can also apply our risk bound to the rare/weak feature model defined in Example 2.2. We have the following corollary.

Corollary 3.8 (Rare/weak feature model)**.** Consider the rare/weak feature model (Example 2.2). Suppose that

$$d \geq C \max\{n^2, \gamma n \sqrt{s \log(n)}\}$$

and $\gamma\sqrt{s} \geq C$ for some large enough absolute constant C. Then when n is large enough, with probability at least $1 - n^{-1}$, the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ has the following risk bound

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\bigg(-\frac{C' n \gamma^4 s^2}{n\gamma^2 s + d}\bigg),$$

where C' is an absolute constant.

By Corollary 3.8, we can see that our bound is tighter by a factor of n in the exponent compared with the risk bound in Chatterji and Long (2020) for the rare/weak feature model. Under the setting where n and γ are fixed constants, our bound can also be compared with the negative result in Jin (2009), which showed that achieving a small population risk is impossible when $s = O(d^2)$. Our result, on the other hand, demonstrates that when $s = \omega(d^2)$, $o(1)$ population risk is achievable.

4 Proof of the Main Results

In this section, we explain how we establish the population risk bound of the maximum margin classifier, and give the proof of Theorem 3.1.

For classification problems, one of the key challenges is that the maximum margin classifier usually does not have an explicit form solution. To overcome this difficulty, Chatterji and Long (2020) utilized the implicit bias results (Lemma 2.3) to get a handle on the relationship between the maximum margin classifier and the training data. More recently, Wang and Thrampoulidis (2020) showed that for isotropic Gaussian mixture models, an explicit form of $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ can be calculated by the equivalence between hard-margin support vector machine and minimum norm least square regression. Notably, it was shown that such an equivalence result holds under the assumptions of Chatterji and Long (2020) and no any additional assumptions are needed. In this paper, we also study the equivalence between classification and regression as a first step. However, our proof works for a more general setting that covers both isotropic and anisotropic sub-Gaussian mixtures, and introduces a novel proof technique based on the polarization identity that leads to a tighter bound. We present this result in Section 4.1.

4.1 Equivalence Between Classification and Regression

Here we establish an equivalence guarantee for the maximum margin classifier and the minimum norm interpolator. We first define the minimum norm interpolator $\widehat{\boldsymbol{\theta}}_{\mathrm{LS}}$ as follows:

$$\widehat{\boldsymbol{\theta}}_{\mathrm{LS}} := \mathrm{argmin}\, \|\boldsymbol{\theta}\|_2^2, \ \ \text{subject to } y_i \cdot \langle \boldsymbol{\theta}, \mathbf{x}_i \rangle = 1, i \in [n].$$

In comparison, we recall that the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ is defined in Lemma 2.3 as

$$\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \mathrm{argmin}\, \|\boldsymbol{\theta}\|_2^2, \ \ \text{subject to } y_i \cdot \langle \boldsymbol{\theta}, \mathbf{x}_i \rangle \geq 1, i \in [n].$$

We can see that the two optimization problems have the same solution when all the training data are support vectors, i.e., all the inequalities become equalities in the constraints. Based on this observation, Muthukumar et al. (2020a); Hsu et al. (2020) have studied the conditions under which the maximum margin classifier $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}$ is identical to the minimum norm interpolator $\widehat{\boldsymbol{\theta}}_{\mathrm{LS}}$. The result is given in the following lemma.

Lemma 4.1 (Hsu et al. (2020))**.** $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \widehat{\boldsymbol{\theta}}_{\mathrm{LS}}$ if and only if

$$\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i > 0$$

for all $i \in [n]$.

According to Lemma 4.1, to study the equivalence between the maximum margin classifier and the minimum norm interpolator, it suffices to derive sufficient conditions such that $\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i$, $i \in [n]$ are strictly positive with high probability. We have the following lemma which summarizes some calculations regarding the quantity $\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i$.

Lemma 4.2. Suppose that

$$\mathrm{tr}(\boldsymbol{\Sigma}) > C\max\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$$

for some absolute constant C. Then with probability at least $1 - O(n^{-2})$,

$$\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i \geq G\Big[1 - C'n\big|\boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i\big|\Big]$$

for all $i \in [n]$, where $G = G(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) > 0$ is a strictly positive factor and $C' > 0$ is an absolute constant.

By Lemma 4.2, we can see that in order to ensure $\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i > 0$, it suffices to establish an upper bound for $|\boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i|$. However, deriving tight upper bounds for this term turns out to be challenging, as a simple application of the Cauchy-Schwarz inequality can lead to a loose bound with an additional $\sqrt{n}$ factor. In the following, we establish a refined bound on the term $|\boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i|$.

Lemma 4.3. Suppose that

$$\mathrm{tr}(\boldsymbol{\Sigma}) > C\max\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F\}$$

for some absolute constant C. Then with probability at least $1 - O(n^{-2})$,

$$\big|\boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i\big| \leq \frac{C'\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \cdot \sqrt{\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma})}$$

for all $i \in [n]$, where $C' > 0$ is an absolute constant.

The proof of Lemma 4.3 is motivated by the observation that the matrix $\mathbf{Q}\mathbf{Q}^\top \in \mathbb{R}^{n\times n}$ is close enough to a scalar multiple of the identity matrix. By concentration arguments, we have with high probability that,

$$\|\mathbf{Q}\mathbf{Q}^\top - \mathrm{tr}(\boldsymbol{\Sigma})\mathbf{I}\|_2 \leq \epsilon_\lambda,$$

where $\epsilon_\lambda := C \cdot \big(n \cdot \|\boldsymbol{\Sigma}\|_2 + \sqrt{n} \cdot \|\boldsymbol{\Sigma}\|_F\big)$ (See Lemma A.1 and its proof for more details). We can then apply the polarization identity $\mathbf{a}^\top\mathbf{M}\mathbf{b} = (1/4)\cdot(\mathbf{a}+\mathbf{b})^\top\mathbf{M}(\mathbf{a}+\mathbf{b}) - (1/4)\cdot(\mathbf{a}-\mathbf{b})^\top\mathbf{M}(\mathbf{a}-\mathbf{b})$

to obtain

$$
\begin{aligned}
\boldsymbol{\mu}^\top \mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i &= \frac{(\mathbf{Q}\boldsymbol{\mu} + \|\mathbf{Q}\boldsymbol{\mu}\|_2 \mathbf{e}_i)^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1} (\mathbf{Q}\boldsymbol{\mu} + \|\mathbf{Q}\boldsymbol{\mu}\|_2 \mathbf{e}_i)}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \\
&\quad - \frac{(\mathbf{Q}\boldsymbol{\mu} - \|\mathbf{Q}\boldsymbol{\mu}\|_2 \mathbf{e}_i)^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1} (\mathbf{Q}\boldsymbol{\mu} - \|\mathbf{Q}\boldsymbol{\mu}\|_2 \mathbf{e}_i)}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \\
&\le \frac{1}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \left[\frac{\|\mathbf{Q}\boldsymbol{\mu} + \|\mathbf{Q}\boldsymbol{\mu}\|_2 \mathbf{e}_i\|_2^2}{\operatorname{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} - \frac{\|\mathbf{Q}\boldsymbol{\mu} - \|\mathbf{Q}\boldsymbol{\mu}\|_2 \mathbf{e}_i\|_2^2}{\operatorname{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda} \right] \\
&= \frac{\|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \epsilon_\lambda + \mathbf{e}_i^\top \mathbf{Q}\boldsymbol{\mu} \cdot \operatorname{tr}(\boldsymbol{\Sigma})}{\operatorname{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2} \\
&\le \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot O\bigg(\frac{\epsilon_\lambda}{\operatorname{tr}(\boldsymbol{\Sigma})^2}\bigg) + |\mathbf{e}_i^\top \mathbf{Q}\boldsymbol{\mu}| \cdot O\bigg(\frac{1}{\operatorname{tr}(\boldsymbol{\Sigma})}\bigg),
\end{aligned} \tag{4.1}
$$

where the first inequality follows by the eigenvalue bound $\|\mathbf{Q}\mathbf{Q}^\top - \operatorname{tr}(\boldsymbol{\Sigma})\mathbf{I}\| \le \epsilon_\lambda$, and the second inequality follows because the lemma assumption implies that $\operatorname{tr}(\boldsymbol{\Sigma}) \ge C\sqrt{n}\epsilon_{\boldsymbol{\lambda}}$ for some large enough absolute constant C. By concentration arguments, with high probability we have $\|\mathbf{Q}\boldsymbol{\mu}\|_2 = O(\sqrt{n}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})$ and $|\mathbf{e}_i^\top \mathbf{Q}\boldsymbol{\mu}| = O(\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})$ (See the proof of Lemma 4.3 in Appendix A.1 for more details). Note that the same upper bound also holds for $-\boldsymbol{\mu}^\top \mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i$ with exactly the same proof. Therefore, plugging these calculations into (4.1) and applying the assumption $\operatorname{tr}(\boldsymbol{\Sigma}) \ge C\sqrt{n}\epsilon_{\boldsymbol{\lambda}}$ completes the proof.

From the above proof sketch of Lemma 4.3, we can see that our proof technique enables us to show that $|\boldsymbol{\mu}^\top \mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i|$ has an upper bound (4.1) whose leading term only depends on $\mathbf{e}_i^\top \mathbf{Q}\boldsymbol{\mu}$. Note that the bound of $\mathbf{e}_i^\top \mathbf{Q}\boldsymbol{\mu}$ is (almost) independent of n because it is only related to the i-th training example. This is a key feature of Lemma 4.3 which leads to a tight bound.

We are now ready to present our main result on the equivalence between maximum margin classifier and the minimum norm interpolator.

Proposition 4.4. Suppose that

$$
\operatorname{tr}(\boldsymbol{\Sigma}) \ge C \max\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}
$$

for some absolute constant C. Then with probability at least $1 - O(n^{-2})$, $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \widehat{\boldsymbol{\theta}}_{\mathrm{LS}}$.

Proof of Proposition 4.4. By the union bound, we have that with probability at least $1 - 2n^{-2}$, the results in Lemma 4.2 and Lemma 4.3 both hold. Therefore, for any $i \in [n]$, we have

$$
\begin{aligned}
\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i &\ge G\Big[1 - c_1 n\big|\boldsymbol{\mu}^\top \mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{e}_i\big|\Big] \\
&\ge G\bigg[1 - \frac{c_2 n\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}}{\operatorname{tr}(\boldsymbol{\Sigma})}\bigg] \\
&\propto \operatorname{tr}(\boldsymbol{\Sigma}) - c_2 n\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}.
\end{aligned}
$$

By the assumption $\operatorname{tr}(\boldsymbol{\Sigma}) \ge Cn\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ for some large enough absolute constant C, we have $\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i > 0$. Finally, applying Lemma 4.1, we conclude that $\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \widehat{\boldsymbol{\theta}}_{\mathrm{LS}}$. $\square$

4.2 Population Risk Bound for the Maximum Margin Classifier

In this subsection, we derive the population risk bound for the maximum margin classifier and provide the proof of Theorem 3.1. We first present the following lemma on the risk bound of linear classifiers for sub-Gaussian mixture models.

Lemma 4.5. There exists an absolute constant C, such that for any $\boldsymbol{\theta} \in \mathbb{R}^d$, the following risk bound holds:

$$R(\boldsymbol{\theta}) \leq \exp\bigg(-\frac{C(\boldsymbol{\theta}^\top \boldsymbol{\mu})^2}{\|\boldsymbol{\theta}\|_{\boldsymbol{\Sigma}}^2}\bigg).$$

A similar result is given in Chatterji and Long (2020) where $\boldsymbol{\Sigma}$ is replaced by $\mathbf{I}$. Our result here depends on the full spectrum of the covariance matrix and is sharper than Chatterji and Long (2020) when $\boldsymbol{\Sigma}$ has decaying eigenvalues.

The proof of Lemma 4.5 is given in Appendix A.1. In addition to this risk bound for general vector $\boldsymbol{\theta}$, we also have the following explicit calculation for $\widehat{\boldsymbol{\theta}}_{\text{SVM}}$ thanks to our analysis in Section 4.1. This is because the minimum norm interpolator $\widehat{\boldsymbol{\theta}}_{\text{LS}}$ has the explicit form $\widehat{\boldsymbol{\theta}}_{\text{LS}} = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}$. Therefore by Proposition 4.4, we also have

$$\widehat{\boldsymbol{\theta}}_{\text{SVM}} = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}.$$

Plugging the above calculation into the risk bound in Lemma 4.5 and utilizing the model definition $\mathbf{X} = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q}$, we are able to show the following risk bound for $\widehat{\boldsymbol{\theta}}_{\text{SVM}}$.

Lemma 4.6. Suppose that

$$\operatorname{tr}(\boldsymbol{\Sigma}) \geq C\max\{n\sqrt{\log(n)}, n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$$

for some absolute constant C. Then with probability at least $1 - O(n^{-2})$,

$$R(\widehat{\boldsymbol{\theta}}_{\text{SVM}}) \leq \exp\bigg\{\frac{-C' \cdot [\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2}{(\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y})^2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2}\bigg\},$$

where C' is an absolute constant.

The proof of Lemma 4.6 is given in Appendix A.1. Lemma 4.6 utilizes the structure of the model to divide the denominator in the exponent into two terms. Motivated by this result, we define

$$\begin{aligned}
I_1 &= [\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2,\\
I_2 &= (\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y})^2 \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2,\\
I_3 &= \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2.
\end{aligned}$$

In the following, we develop a lower bound for I_1 and upper bounds for I_2 and I_3 respectively. The following lemma summarizes the bounds.

Lemma 4.7. Suppose that

$$\operatorname{tr}(\boldsymbol{\Sigma}) \geq C\max\{n, n\|\boldsymbol{\Sigma}\|_2, \sqrt{n}\|\boldsymbol{\Sigma}\|_F, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$$

and $\|\boldsymbol{\mu}\|_2 \geq C\|\boldsymbol{\Sigma}\|_2$ for some absolute constant C. Then when n is large enough, with probability at least $1 - O(n^{-2})$,

$$\begin{aligned}
&I_1 \geq C'^{-1} H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot n^2 \cdot \|\boldsymbol{\mu}\|_2^4, \\
&I_2 \leq C' H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot n^2 \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2, \\
&I_3 \leq C' H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2),
\end{aligned}$$

where $H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) > 0$ is a strictly positive coefficient, and $C' > 0$ is an absolute constant.

The proof of Lemma 4.7 is given in Appendix A.1. To illustrate the key idea in the proof of Lemma 4.7, we take I_3 as an example. Based on our model in Section 2, we have $\mathbf{Q} = \mathbf{Z}\boldsymbol{\Lambda}^{1/2}\mathbf{V}^\top$, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is a random matrix with independent sub-Gaussian entries, and $\boldsymbol{\Lambda}$, $\mathbf{V}$ are defined based on the eigenvalue decomposition $\boldsymbol{\Sigma} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top$. By some linear algebra calculation (see the proof of Lemma 4.7 in Appendix A.1 for more details), we have

$$I_3 = \mathbf{a}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1} \mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{a}, \tag{4.2}$$

where $\|\mathbf{a}\|_2^2 = O(D^{-2}n)$ with

$$D = \mathbf{y}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{y} \cdot (\|\boldsymbol{\mu}\|_2^2 - \boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu}) + (1 + \mathbf{y}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu})^2.$$

The key observation here is that while the term D above has a very complicated form, it is not necessary to bound it. This is because D^{-2} is a common term that appears in all I_1, I_2 and I_3 and therefore can be canceled out when calculating the ratio $I_1/(I_2 + I_3)$. With the calculation in (4.2), we are able to invoke the following eigenvalue concentration inequalities (see Lemma A.1 and Lemma A.5 for more details) to give upper and lower bounds regarding the matrices $\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top$ and $\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top$ respectively:

$$\begin{aligned}
&\big\|\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top - \mathrm{tr}(\boldsymbol{\Sigma}) \cdot \mathbf{I}\big\|_2 \leq c_1 \cdot \big(n\|\boldsymbol{\Sigma}\|_2 + \sqrt{n}\|\boldsymbol{\Sigma}\|_F\big), \\
&\big\|\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top - \|\boldsymbol{\Sigma}\|_F^2 \cdot \mathbf{I}\big\|_2 \leq c_1 \cdot \big(n\|\boldsymbol{\Sigma}\|_2^2 + \sqrt{n}\|\boldsymbol{\Sigma}^2\|_F\big),
\end{aligned}$$

where c_1 is an absolute constant. Plugging the above inequalities and the bound $\|\mathbf{a}\|_2^2 = O(D^{-2}n)$ into (4.2), we obtain with some calculation that

$$I_3 \leq c_2 H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2)$$

with $H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) = [D \cdot \mathrm{tr}(\boldsymbol{\Sigma})]^{-2}$, where c_2 is an absolute constant. This gives the bound of I_3 in Lemma 4.7.

Lemma 4.7 is significant in three-fold. First of all, the result does not have an explicit dependency on d, which makes it applicable to infinite dimensional data. Second, Lemma 4.7 gives bounds with great simplicity, and shows that the three bounds share a same strictly positive factor $H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})$, which can be canceled out since our final goal is to bound the ratio $I_1/(I_2 + I_3)$. Lastly, Lemma 4.7 reveals the fact that the risk bound only depends on $\|\boldsymbol{\Sigma}\|_F$ and $\|\boldsymbol{\Sigma}\|_2$, which can be small even though the assumption requires $\mathrm{tr}(\boldsymbol{\Sigma})$ to be large.

We are now ready to present the proof of Theorem 3.1.

Proof of Theorem 3.1. Clearly, under the assumptions of Theorem 3.1, the conditions in Lemma 4.6

and Lemma 4.7 are both satisfied. By Lemma 4.7, we have

$$\frac{[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2}{(\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y})^2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2} \geq c_1 \cdot \frac{n^2\|\boldsymbol{\mu}\|_2^4}{n^2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + n\cdot\|\boldsymbol{\Sigma}\|_F^2 + n^2\cdot\|\boldsymbol{\Sigma}\|_2^2},$$

where c_1 is an absolute constant. Therefore by Lemma 4.6 we have

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\left(\frac{-c_2 n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2}\right)$$

for some absolute constant c_2. Note that by union bound, the above inequality holds with probability at least $1 - O(n^{-2}) \geq 1 - n^{-1}$ when n is large enough. This completes the proof. $\square$

5 Conclusion and Future Work

We have studied the benign overfitting phenomenon for sub-Gaussian mxiture models, and established a population risk bound for the maximum margin classifier. Our population risk bound is general and covers both the isotropic and anisotropic settings. When reduced to the isotropic setting, our bound is tighter than existing results. We have also studied a class of non-isotropic models which can be benign even for infinite-dimensional data.

An interesting future work direction is to study the relation between the dimension and the population risk and verify the double descent phenomenon. Studying benign overfitting for more complicated learning models such as neural networks is another important future work direction.

A Completing the Proof of Theorem 3.1

In Section 4 we give the proof of Theorem 3.1 based on several technical lemmas. Here we present the complete proof of these lemmas.

A.1 Proof of Lemmas in Section 4

We denote $\boldsymbol{\nu} = \mathbf{Q}\boldsymbol{\mu}$ and $\mathbf{A} = \mathbf{Q}\mathbf{Q}^\top$. Based on these notations, in the following we present several basic lemmas that are used in our proof. We have the following lemma which gives concentration inequalities for the the eigenvalues of $\mathbf{A}$.

Lemma A.1. With probability at least $1 - n^{-2}$,

$$\big\|\mathbf{A} - \mathrm{tr}(\boldsymbol{\Sigma})\cdot\mathbf{I}\big\|_2 \leq \epsilon_\lambda := C\sigma_u^2\big(n\cdot\|\boldsymbol{\Sigma}\|_2 + \sqrt{n}\cdot\|\boldsymbol{\Sigma}\|_F\big),$$

where C is an absolute constant.

The following lemma presents some calculations on the quantity $\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}$. It utilizes a result introduced in Wang and Thrampoulidis (2020), which is based on the application of the Sherman–Morrison–Woodbury formula.

Lemma A.2. The following calculation of $\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}$ holds:

$$\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1} = D^{-1}[(1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\cdot\mathbf{y}^\top\mathbf{A}^{-1} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}],$$

where $D = \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \cdot (\|\boldsymbol{\mu}\|_2^2 - \boldsymbol{\nu}^\top \mathbf{A}^{-1}\boldsymbol{\nu}) + (1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu})^2 > 0$.

Motivated by Lemma A.2, we estimate the orders of the terms $\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y}$, $\boldsymbol{\nu}^\top \mathbf{A}^{-1}\boldsymbol{\nu}$, and $\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}$. The results are given in the following lemma.

Lemma A.3. Let ϵ_λ be defined in Lemma A.1, and suppose that $\mathrm{tr}(\boldsymbol{\Sigma}) > \epsilon_\lambda$. Then with probability at least $1 - O(n^{-2})$, the following inequalities hold:

$$
\begin{aligned}
&\frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda} \le \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \le \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda},\\
&\frac{n - C\sqrt{n\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \le \boldsymbol{\nu}^\top \mathbf{A}^{-1}\boldsymbol{\nu} \le \frac{n + C\sqrt{n\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2,\\
&|\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}| \le \frac{Cn}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}},
\end{aligned}
$$

where C is an absolute constant.

A.1.1 Proof of Lemma 4.2

Here we present the proof of Lemma 4.2. We first introduce the following lemma, which gives a lower bound on the quantity $\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i$.

Lemma A.4. Let ϵ_λ be defined in Lemma A.1, and suppose that $\mathrm{tr}(\boldsymbol{\Sigma}) > \epsilon_\lambda$. Then with probability at least $1 - O(n^{-2})$,

$$
\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i \ge \frac{\mathrm{tr}(\boldsymbol{\Sigma}) - \sqrt{n}\epsilon_\lambda}{\mathrm{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2}
$$

for all $i \in [n]$.

We are now ready to give our proof of Lemma 4.2.

Proof of Lemma 4.2. By Lemma A.2, we have

$$
\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i = D^{-1}[(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i - \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \cdot \boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i].
$$

Plugging in the inequalities in Lemmas A.3 and A.4, we have that as long as $\mathrm{tr}(\boldsymbol{\Sigma}) > c_1 \max\{n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}, \epsilon_\lambda\}$ for some large enough constant c_1, $\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu} \le 1/2$ and therefore

$$
\begin{aligned}
\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i &= D^{-1}[(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i - \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \cdot \boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i]\\
&\ge D^{-1} \cdot \left[\frac{1}{2} \cdot \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i - \frac{c_2 n}{\mathrm{tr}(\boldsymbol{\Sigma})} \cdot |\boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i|\right],
\end{aligned} \tag{A.1}
$$

where c_2 is an absolute constant. By Lemma A.4, we can see that as long as $\mathrm{tr}(\boldsymbol{\Sigma}) \ge c_3\sqrt{n}\epsilon_\lambda$ for some large enough absolute constant c_3, we have

$$
\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i \ge \frac{\mathrm{tr}(\boldsymbol{\Sigma}) - \sqrt{n}\epsilon_\lambda}{\mathrm{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2} \ge \frac{1}{2\,\mathrm{tr}(\boldsymbol{\Sigma})}.
$$

Plugging the bound above into (A.1), we obtain

$$\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{e}_i y_i \geq \frac{1}{4D\,\mathrm{tr}(\boldsymbol{\Sigma})} \cdot [1 - c_4 n \cdot |\boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i|].$$

Since $D > 0$, we see that $G(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) := [4D\,\mathrm{tr}(\boldsymbol{\Sigma})]^{-1} > 0$. This completes the proof. $\square$

A.1.2 Proof of Lemma 4.3

Here we give the detailed proof of Lemma 4.3 to backup the proof sketch presented in Section 4.1. The proof is based on the polarization identity.

Proof of Lemma 4.3. We have the following calculation,

$$\begin{aligned}
\boldsymbol{\mu}^\top\mathbf{Q}^\top\mathbf{A}^{-1}\mathbf{e}_i y_i &= \frac{1}{\|\mathbf{Q}\boldsymbol{\mu}\|_2} \cdot (\mathbf{Q}\boldsymbol{\mu})^\top \mathbf{A}^{-1}(\|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i) \\
&= \frac{1}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \cdot (\mathbf{Q}\boldsymbol{\mu} + \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i)^\top \mathbf{A}^{-1}(\mathbf{Q}\boldsymbol{\mu} + \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i) \\
&\quad - \frac{1}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \cdot (\mathbf{Q}\boldsymbol{\mu} - \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i)^\top \mathbf{A}^{-1}(\mathbf{Q}\boldsymbol{\mu} - \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i) \\
&\leq \frac{1}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \cdot \left[\frac{\|\mathbf{Q}\boldsymbol{\mu} + \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i\|_2^2}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} - \frac{\|\mathbf{Q}\boldsymbol{\mu} - \|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i y_i\|_2^2}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda}\right] \\
&= \frac{1}{4\|\mathbf{Q}\boldsymbol{\mu}\|_2} \cdot \left[\frac{2\|\mathbf{Q}\boldsymbol{\mu}\|_2^2 + 2y_i\|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i^\top\mathbf{Q}\boldsymbol{\mu}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} - \frac{2\|\mathbf{Q}\boldsymbol{\mu}\|_2^2 - 2y_i\|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i^\top\mathbf{Q}\boldsymbol{\mu}}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda}\right] \\
&= \frac{1}{2\|\mathbf{Q}\boldsymbol{\mu}\|_2} \cdot \frac{2\|\mathbf{Q}\boldsymbol{\mu}\|_2^2 \cdot \epsilon_\lambda + 2y_i\|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \mathbf{e}_i^\top\mathbf{Q}\boldsymbol{\mu} \cdot \mathrm{tr}(\boldsymbol{\Sigma})}{\mathrm{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2} \\
&= \frac{\|\mathbf{Q}\boldsymbol{\mu}\|_2 \cdot \epsilon_\lambda + y_i\mathbf{e}_i^\top\mathbf{Q}\boldsymbol{\mu} \cdot \mathrm{tr}(\boldsymbol{\Sigma})}{\mathrm{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2}, \qquad (A.2)
\end{aligned}$$

where the first equality holds due to the polarization identity $\mathbf{a}^\top\mathbf{M}\mathbf{b} = 1/4(\mathbf{a}+\mathbf{b})^\top\mathbf{M}(\mathbf{a}+\mathbf{b}) - 1/4(\mathbf{a}-\mathbf{b})^\top\mathbf{M}(\mathbf{a}-\mathbf{b})$, and the first inequality follows by Lemma A.1. Based on our model assumption, we can denote $\mathbf{Q} = \mathbf{Z}\boldsymbol{\Lambda}^{1/2}\mathbf{V}^\top$, where the entries of $\mathbf{Z}$ are independent sub-Gaussian random variables with $\|\mathbf{Z}_{ij}\|_{\psi_2} \leq \sigma_u$ for all $i \in [n]$ and $j \in [p]$. Denote $\widetilde{\boldsymbol{\mu}} = \Lambda^{1/2}\mathbf{V}^\top\boldsymbol{\mu}$. Then with the same proof as in Lemma A.3, we have

$$\|\mathbf{Q}\boldsymbol{\mu}\|_2^2 = \|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2 \leq 2n\|\widetilde{\boldsymbol{\mu}}\|_2^2 = 2n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2$$

when n is large enough. Moreover, we also have

$$\|y_i\mathbf{e}_i^\top\mathbf{Q}\boldsymbol{\mu}\|_{\psi_2} = \left\|\sum_{j=1}^p \mathbf{Z}_{ij}\widetilde{\mu}_j\right\|_{\psi_2} \leq \|\widetilde{\boldsymbol{\mu}}\|_2 \cdot \sigma_u.$$

Therefore by Hoeffding's inequality, with probability at least $1 - n^{-1}$, we have

$$|y_i\mathbf{e}_i^\top\mathbf{Q}\boldsymbol{\mu}| \leq c_1\|\widetilde{\boldsymbol{\mu}}\|_2 \cdot \sqrt{\log(n)} = c_1\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \cdot \sqrt{\log(n)},$$

where c_1 is an absolute constant. Therefore we have

$$\boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i \leq \frac{\sqrt{2n}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \cdot \epsilon_\lambda + c_2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\sqrt{\log(n)} \cdot \mathrm{tr}(\boldsymbol{\Sigma})}{\mathrm{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2}.$$

With the exact same proof, we also have

$$-\boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i \leq \frac{\sqrt{2n}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \cdot \epsilon_\lambda + c_2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\sqrt{\log(n)} \cdot \mathrm{tr}(\boldsymbol{\Sigma})}{\mathrm{tr}(\boldsymbol{\Sigma})^2 - \epsilon_\lambda^2}.$$

Therefore by the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) > C\sqrt{n}\epsilon_\lambda$ for some large enough absolute constant C, we have

$$|\boldsymbol{\nu}^\top \mathbf{A}^{-1}\mathbf{e}_i| \leq \frac{c_3\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \cdot \sqrt{\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma})}$$

for some absolute constant c_3. This completes the proof. □

A.1.3 Proof of Lemma 4.5

Here we give the detailed proof of Lemma 4.5, which is based on the one-side sub-Gaussian tail bound.

Proof of Lemma 4.5. By definition, we have

$$R(\boldsymbol{\theta}) = \mathbb{P}(y \cdot \boldsymbol{\theta}^\top \mathbf{x} < 0) = \mathbb{P}[y \cdot \boldsymbol{\theta}^\top (y \cdot \boldsymbol{\mu} + \mathbf{q}) < 0] = \mathbb{P}[\boldsymbol{\theta}^\top \boldsymbol{\mu} < y \cdot \boldsymbol{\theta}^\top \mathbf{q}] = \mathbb{P}[\boldsymbol{\theta}^\top \boldsymbol{\mu} < y \cdot \boldsymbol{\theta}^\top \mathbf{V}\boldsymbol{\Lambda}^{1/2}\mathbf{u}],$$

where in the second and last equations we plug in the definitions of $\mathbf{x}$ and $\mathbf{q}$ according to our data generation procedure described in Section 2. Note that $\mathbf{u}$ has independent, σ_u-sub-Gaussian entries. Therefore we have

$$\|\boldsymbol{\theta}^\top \mathbf{V}\boldsymbol{\Lambda}^{1/2}\mathbf{u}\|_{\psi_2} \leq c_1\|\boldsymbol{\theta}^\top \mathbf{V}\boldsymbol{\Lambda}^{1/2}\|_2 = c_1\sqrt{\boldsymbol{\theta}^\top \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top \boldsymbol{\theta}} = c_1\sqrt{\boldsymbol{\theta}^\top \boldsymbol{\Sigma}\boldsymbol{\theta}}.$$

Applying the one-side sub-Gaussian tail bound (e.g., Theorem A.2 in Chatterji and Long (2020)) completes the proof. □

A.1.4 Proof of Lemma 4.6

The proof of Lemma 4.6 is given as follows, where we utilize Proposition 4.4 and Lemma 4.5 to derive the desired bound.

Proof of Lemma 4.6. By Proposition 4.4, we have

$$\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \widehat{\boldsymbol{\theta}}_{\mathrm{LS}} = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}.$$

Plugging it into the risk bound in Lemma 4.5, we obtain

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\bigg\{-\frac{C[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2}{\|\mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2}\bigg\}.$$

Note that based on our model, we have $\mathbf{X} = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q}$, and

$$\begin{aligned}\|\mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2 &= \|(\mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q})^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2 \\ &\le 2\|\boldsymbol{\mu}\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2 + 2\|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2 \\ &= 2(\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y})^2\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + 2\|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2.\end{aligned}$$

Therefore we have

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \le \exp\left\{\frac{-(C/2)\cdot[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2}{(\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y})^2\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2}\right\}.$$

This completes the proof. □

A.1.5 Proof of Lemma 4.7

In this subsection we present the proof of Lemma 4.7. We first give the following lemma, which follows by exactly the same proof as Lemma A.1.

Lemma A.5. Suppose that $\mathbf{Z} \in \mathbb{R}^{n\times d}$ is a random matrix with i.i.d. sub-Gaussian entries with sub-Gaussian norm σ_u. Then with probability at least $1 - O(n^{-2})$,

$$\big\|\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top - \|\boldsymbol{\Sigma}\|_F^2\cdot\mathbf{I}\big\|_2 \le \epsilon_\lambda' := C\sigma_u^2\big(n\cdot\|\boldsymbol{\Sigma}\|_2^2 + \sqrt{n}\cdot\|\boldsymbol{\Sigma}^2\|_F\big),$$

where C is an absolute constant.

Based on Lemma A.5, we can give the proof of Lemma 4.7 as follows.

Proof of Lemma 4.7. We first derive the lower bound for I_1. By Lemma A.2 and the model definition $\mathbf{X} = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q}$, we have

$$\begin{aligned}\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu} &= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}](\mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q})\boldsymbol{\mu} \\ &= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}](\mathbf{y}\cdot\|\boldsymbol{\mu}\|_2^2 + \mathbf{Q}\boldsymbol{\mu}) \\ &= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}\mathbf{y}]\cdot\|\boldsymbol{\mu}\|_2^2 \\ &\quad + D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu}]) \\ &= D^{-1}\cdot[(\|\boldsymbol{\mu}\|_2^2 - \boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} + (1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}],\end{aligned} \tag{A.3}$$

where the third equality follows by the notation $\boldsymbol{\nu} = \mathbf{Q}\boldsymbol{\mu}$. By Lemma A.3 and the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \ge C\max\{\epsilon_\lambda, n\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$ for some large enough constant C, when n is large enough we have

$$\begin{aligned}&|\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}| \le \frac{c_1 n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \le \frac{2c_1 n}{\mathrm{tr}(\boldsymbol{\Sigma})}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \le 1, \\ &0 \le \boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu} \le \frac{n + c_2\sqrt{n\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \le \frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \le \frac{2n\|\boldsymbol{\Sigma}\|_2}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_2^2 \le \frac{1}{2}\cdot\|\boldsymbol{\mu}\|_2^2, \\ &\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} \ge \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda} \ge \frac{n}{2\,\mathrm{tr}(\boldsymbol{\Sigma})},\end{aligned}$$

where c_1, c_2 are absolute constants. Plugging the bounds above into (A.3), we obtain

$$
\begin{aligned}
|\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}| &\geq D^{-1}\cdot\bigg(\frac{1}{2}\cdot\|\boldsymbol{\mu}\|_2^2\cdot\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} - 2\cdot|\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}|\bigg)\\
&\geq D^{-1}\cdot\bigg[\frac{n}{4\,\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_2^2 - \frac{4n}{\mathrm{tr}(\boldsymbol{\Sigma})}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\bigg]\\
&\geq D^{-1}\cdot\frac{n}{4\,\mathrm{tr}(\boldsymbol{\Sigma})}\cdot(\|\boldsymbol{\mu}\|_2^2 - 16\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})\\
&\geq D^{-1}\cdot\frac{n}{8\,\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_2^2,
\end{aligned}
$$

where the last inequality follows by the assumption that $\|\boldsymbol{\mu}\|_2^2 \geq C\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ for some large enough absolute constant C. Therefore we have

$$
[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2 \geq D^{-2}\cdot\frac{n^2}{64[\mathrm{tr}(\boldsymbol{\Sigma})]^2}\cdot\|\boldsymbol{\mu}\|_2^4 = \frac{H(\boldsymbol{\mu},\mathbf{Q},\mathbf{y},\boldsymbol{\Sigma})}{64}\cdot n^2\|\boldsymbol{\mu}\|_2^4,
$$

where we define

$$
H(\boldsymbol{\mu},\mathbf{Q},\mathbf{y},\boldsymbol{\Sigma}) := [D\cdot\mathrm{tr}(\boldsymbol{\Sigma})]^{-2} > 0.
$$

This completes the proof of the lower bound of I_1.

For I_2, by Lemma A.2 we have

$$
\begin{aligned}
\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y} &= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}\mathbf{y}]\\
&= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}\mathbf{y}]\\
&= D^{-1}\cdot\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\\
&\leq D^{-1}\cdot\frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\\
&\leq 2D^{-1}\cdot\frac{n}{\mathrm{tr}(\boldsymbol{\Sigma})},
\end{aligned}
$$

where the first inequality follows by Lemma A.3, and the second inequality follows by the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \geq C\epsilon_\lambda$ for some large enough constant C. Therefore we have

$$
I_2 = (\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y})^2\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \leq 4D^{-2}\cdot\frac{n^2\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2}{[\mathrm{tr}(\boldsymbol{\Sigma})]^2} = 4H(\boldsymbol{\mu},\mathbf{Q},\mathbf{y},\boldsymbol{\Sigma})\cdot n^2\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2,
$$

where we use the definition $H(\boldsymbol{\mu},\mathbf{Q},\mathbf{y},\boldsymbol{\Sigma}) = [D\cdot\mathrm{tr}(\boldsymbol{\Sigma})]^{-2}$. This proves the upper bound of I_2.

For I_3, by our calculation in Lemma A.2, we have

$$
\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1} = D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}]^\top\mathbf{A}^{-1}.
$$

Denote $\mathbf{a} = D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\cdot\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}]$. Then

$$
\begin{aligned}
I_3 &= \mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{Q}\boldsymbol{\Sigma}\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\\
&= \mathbf{a}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\Sigma}\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{a}\\
&= \mathbf{a}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{a}, \qquad (A.4)
\end{aligned}
$$

where we plug in $\boldsymbol{\Sigma} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top$ and $\mathbf{Q} = \mathbf{Z}\boldsymbol{\Lambda}^{1/2}\mathbf{V}^\top$ for $\mathbf{Z}$ with independent sub-Gaussian entries. By Lemma A.1, Lemma A.5 and (A.4), when $\mathrm{tr}(\boldsymbol{\Sigma}) \geq \epsilon_\lambda$ we have

$$\begin{aligned} I_3 &= \mathbf{a}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{a} \\ &\leq \mathbf{a}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-2}\mathbf{a} \cdot \left[\|\boldsymbol{\Sigma}\|_F^2 + \epsilon'_\lambda\right] \\ &\leq \|\mathbf{a}\|_2^2 \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 + \epsilon'_\lambda}{[\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda]^2}. \end{aligned} \tag{A.5}$$

Here the first inequality follows by Lemma A.5, and the second inequality follows by Lemma A.1. By definition, we have

$$\begin{aligned} \|\mathbf{a}\|_2^2 &= \|D^{-1}(1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}\|_2^2 \\ &\leq 2D^{-2}(1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})^2\|\mathbf{y}\|_2^2 + 2D^{-2}(\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y})^2 \cdot \|\mathbf{Q}\boldsymbol{\mu}\|_2^2. \end{aligned}$$

Then with the same proof as in Lemma A.3, when n is sufficiently large, with probability at least $1 - O(n^{-2})$ we have

$$\|\mathbf{Q}\boldsymbol{\mu}\|_2^2 \leq 2n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2.$$

Therefore we have

$$\begin{aligned} \|\mathbf{a}\|_2^2 &\leq 2D^{-2}(1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})^2\|\mathbf{y}\|_2^2 + 2D^{-2}(\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y})^2 \cdot \|\mathbf{Q}\boldsymbol{\mu}\|_2^2 \\ &\leq 2D^{-2}(1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})^2 \cdot n + 4D^{-2}(\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y})^2 \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2. \end{aligned} \tag{A.6}$$

Moreover, by Lemma A.3 and the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \geq C\max\{\epsilon_\lambda, n, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$ for some large enough constant C, we have

$$\begin{aligned} |\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}| &\leq \frac{c_3 n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \leq \sqrt{2} - 1, \\ \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} &\leq \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} \leq \frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})}, \end{aligned}$$

where c_3 is an absolute constant. Plugging the above bounds into (A.6), we obtain

$$\begin{aligned} \|\mathbf{a}\|_2^2 &\leq 2D^{-2}(1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})^2 \cdot n + 4D^{-2}(\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y})^2 \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \\ &\leq 4D^{-2} \cdot n + 8D^{-2} \cdot n \cdot \left[\frac{n}{\mathrm{tr}(\boldsymbol{\Sigma})} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\right]^2 \\ &\leq 5D^{-2} \cdot n, \end{aligned}$$

where the last inequality utilizes the assumption $\mathrm{tr}(\boldsymbol{\Sigma}) \geq Cn\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ for some large enough constant C again. Further plugging this bound into (A.5), we obtain

$$\begin{aligned} I_3 &\leq \|\mathbf{a}\|_2^2 \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 + \epsilon'_\lambda}{[\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda]^2} \leq 5D^{-2}n \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 + \epsilon'_\lambda}{[\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda]^2} \\ &\leq c_4 D^{-2} \cdot \frac{n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2 + n^{3/2} \cdot \|\boldsymbol{\Sigma}^2\|_F}{[\mathrm{tr}(\boldsymbol{\Sigma})]^2}, \end{aligned} \tag{A.7}$$

where c_4 is an absolute constant. Note that we have

$$n^{3/2} \cdot \|\boldsymbol{\Sigma}^2\|_F \leq n \cdot \|\boldsymbol{\Sigma}\|_F \cdot (\sqrt{n} \cdot \|\boldsymbol{\Sigma}\|_2) \leq n \cdot (\|\boldsymbol{\Sigma}\|_F^2 + n \cdot \|\boldsymbol{\Sigma}\|_2^2)/2.$$

Plugging this bound into (A.7), we have

$$I_3 \leq c_5 D^{-2} \cdot \frac{n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2}{[\mathrm{tr}(\boldsymbol{\Sigma})]^2} = c_5 H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2),$$

where we use the definition $H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) = [D \cdot \mathrm{tr}(\boldsymbol{\Sigma})]^{-2}$, and c_4 is an absolute constant. This finishes the proof of the upper bound of I_3. □

A.2 Proof of Lemmas in Appendix A.1

Here we present the proofs of lemmas we used in Appendix A.1.

A.2.1 Proof of Lemma A.1

The proof of Lemma A.1 is motivated by the analysis given in Bartlett et al. (2020). However here in Lemma A.1 we give a slightly tighter bound. The proof is as follows.

Proof of Lemma A.1. Let $\mathcal{N}$ be a 1/4-net on the unit sphere s^{n-1}. Then by Lemma 5.2 in Vershynin (2010), we have $|\mathcal{N}| \leq 9^n$. Denote $\mathbf{z}_j = \lambda_j^{-1/2} \mathbf{Q} \mathbf{v}_j \in \mathbb{R}^n$. Then by definition, for any fixed unit vector $\widehat{\mathbf{a}} \in \mathcal{N}$ we have $\widehat{\mathbf{a}}^\top \mathbf{A} \widehat{\mathbf{a}} = \mathbf{Q}\mathbf{Q}^\top = \widehat{\mathbf{a}}^\top \sum_{j=1}^p \lambda_j \mathbf{z}_j \mathbf{z}_j^\top \widehat{\mathbf{a}} = \sum_{j=1}^p \lambda_j (\widehat{\mathbf{a}}^\top \mathbf{z}_j)^2$. By Lemma 5.9 in Vershynin (2010), there exists an absolute constant c_1 such that $\|\widehat{\mathbf{a}}^\top \mathbf{z}_j\|_{\psi_2} \leq c_1 \sigma_u$. Therefore by Lemma 21 and Corollary 23 in Bartlett et al. (2020), for any $t > 0$, with probability at least $1 - 2\exp(-t)$ we have

$$\big|\widehat{\mathbf{a}}^\top \mathbf{A} \widehat{\mathbf{a}} - \mathrm{tr}(\boldsymbol{\Sigma})\big| \leq c_2 \sigma_u^2 \max\big(t \cdot \|\boldsymbol{\Sigma}\|_2, \sqrt{t} \cdot \|\boldsymbol{\Sigma}\|_F\big).$$

Applying an union bound over all $\widehat{\mathbf{a}} \in \mathcal{N}$, we have that with probability at least $1 - 2 \cdot 9^n \exp(-t)$,

$$\big|\widehat{\mathbf{a}}^\top \mathbf{A} \widehat{\mathbf{a}} - \mathrm{tr}(\boldsymbol{\Sigma})\big| \leq c_2 \sigma_u^2 \max\big(t \cdot \|\boldsymbol{\Sigma}\|_2, \sqrt{t} \cdot \|\boldsymbol{\Sigma}\|_F\big)$$

for all $\widehat{\mathbf{a}} \in \mathcal{N}$. Therefore by Lemma 25 in Bartlett et al. (2020), with probability at least $1 - 2 \cdot 9^n \exp(-t)$, we have

$$\big\|\mathbf{A} - \mathrm{tr}(\boldsymbol{\Sigma})\mathbf{I}\big\|_2 \leq c_3 \sigma_u^2 \big(t \cdot \|\boldsymbol{\Sigma}\|_2 + \sqrt{t} \cdot \|\boldsymbol{\Sigma}\|_F\big),$$

where c_3 is an absolute constant. Setting $t = c_4 n$ for some large enough constant c_4, we have that with probability at least $1 - n^{-2}$,

$$\big\|\mathbf{A} - \mathrm{tr}(\boldsymbol{\Sigma})\mathbf{I}\big\|_2 \leq c_5 \sigma_u^2 \big(n \cdot \|\boldsymbol{\Sigma}\|_2 + \sqrt{n} \cdot \|\boldsymbol{\Sigma}\|_F\big),$$

where c_5 is an absolute constant. This completes the proof. □

A.2.2 Proof of Lemma A.2

Here we present the proof of Lemma A.2. Our proof utilizes a key lemma by Wang and Thrampoulidis (2020), and gives further simplifications of the result.

Proof of Lemma A.2. Denote $s = \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y}$, $t = \boldsymbol{\nu}^\top \mathbf{A}^{-1}\boldsymbol{\nu}$, $h = \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}$. Then we have $D = \|\boldsymbol{\mu}\|_2^2 s - st + (h+1)^2$. By Lemma 3 in Wang and Thrampoulidis (2020), we have

$$\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1} = \mathbf{y}^\top\mathbf{A}^{-1} - D^{-1}\cdot[\|\boldsymbol{\mu}\|_2^2 s + h^2 + h - st]\cdot \mathbf{y}^\top\mathbf{A}^{-1} - D^{-1}s\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}.$$

Rearranging terms, we obtain

$$\begin{aligned}
\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1} &= \left[1 - \frac{\|\boldsymbol{\mu}\|_2^2 s + h^2 + h - st}{\|\boldsymbol{\mu}\|_2^2 s - st + (h+1)^2}\right]\cdot \mathbf{y}^\top\mathbf{A}^{-1} - D^{-1}s\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1} \\
&= \frac{h+1}{\|\boldsymbol{\mu}\|_2^2 s - st + (h+1)^2}\cdot \mathbf{y}^\top\mathbf{A}^{-1} - D^{-1}s\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1} \\
&= D^{-1}[(h+1)\mathbf{y}^\top\mathbf{A}^{-1} - s\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}].
\end{aligned}$$

At last, by the definition of D, we have

$$\begin{aligned}
D &= \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot(\|\boldsymbol{\mu}\|_2^2 - \boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu}) + (1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})^2 \\
&\geq (1 + \mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})^2,
\end{aligned}$$

where we utilize the fact that $\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} \geq 0$ and $\|\boldsymbol{\mu}\|_2^2 \geq \boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu}$. Since $\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu} \neq 1$ with probability 1, we see that $D > 0$ almost surely. This completes the proof. □

A.2.3 Proof of Lemma A.3

The proof of Lemma A.3 is based on the application of eigenvalue concentration results in Lemma A.1. We present the details as follows.

Proof of Lemma A.3. The bounds on $\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}$ are directly derived from Lemma A.1 and the fact that $\|\mathbf{y}\|_2^2 = n$. To derive the bounds for $\boldsymbol{\nu}\mathbf{A}^{-1}\boldsymbol{\nu}$, we note that by definition, $\boldsymbol{\nu} = \mathbf{Q}\boldsymbol{\mu}$ and

$$\boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu} = \boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu}.$$

Denote $\mathbf{z}_i = \lambda_i^{-1/2}\mathbf{Q}\mathbf{v}_i \in \mathbb{R}^n$, $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_p] \in \mathbb{R}^{n\times p}$, and $\widetilde{\boldsymbol{\mu}} = \Lambda^{1/2}\mathbf{V}^\top\boldsymbol{\mu}$. Then $\mathbf{Q} = \mathbf{Z}\boldsymbol{\Lambda}^{1/2}\mathbf{V}^\top$, $\mathbf{Q}\boldsymbol{\mu} = \mathbf{Z}\widetilde{\boldsymbol{\mu}}$, and

$$\begin{aligned}
\boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu} &= \boldsymbol{\mu}^\top\mathbf{V}\boldsymbol{\Lambda}^{1/2}\mathbf{Z}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{Z}\boldsymbol{\Lambda}^{1/2}\mathbf{V}^\top\boldsymbol{\mu} \\
&= \widetilde{\boldsymbol{\mu}}^\top\mathbf{Z}^\top(\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{Z}\widetilde{\boldsymbol{\mu}} \\
&\leq \frac{\|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}.
\end{aligned}$$

Similarly, we have

$$\boldsymbol{\mu}^\top\mathbf{Q}^\top(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu} \geq \frac{\|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda}.$$

We now proceed to give upper and lower bounds for the term $\|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2 = \sum_{i=1}^n (\sum_{j=1}^p \mathbf{Z}_{ij}\widetilde{\mu}_j)^2$. Note that by definition, $\mathbf{Z}_{ij}$ for $i \in [n]$ and $j \in [p]$ are independent sub-Gaussian vectors with $\|\mathbf{Z}_{ij}\|_{\psi_2} \leq \sigma_u$. By Lemma 5.9 in Vershynin (2010), we have

$$\left\| \sum_{j=1}^p \mathbf{Z}_{ij}\widetilde{\mu}_j \right\|_{\psi_2} \leq c_1 \|\widetilde{\boldsymbol{\mu}}\|_2 \cdot \sigma_u,$$

where c_1 is an absolute constant. Therefore by Lemma 5.14 in Vershynin (2010), we have

$$\left\| \left(\sum_{j=1}^p \mathbf{Z}_{ij}\widetilde{\mu}_j \right)^2 - \|\widetilde{\boldsymbol{\mu}}\|_2^2 \right\|_{\psi_1} \leq c_2 \|\widetilde{\boldsymbol{\mu}}_j\|_2^2,$$

where we merge σ_u into the absolute constant c_2. By Bernstein's inequality, with probability at least $1 - n^{-2}$,

$$\big| \|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2 - \mathbb{E}\|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2 \big| \leq c_3 \|\widetilde{\boldsymbol{\mu}}\|_2^2 \cdot \sqrt{n \log(n)},$$

where c_3 is an absolute constant. Therefore we have

$$n\|\widetilde{\boldsymbol{\mu}}\|_2^2 - c_3\|\widetilde{\boldsymbol{\mu}}\|_2^2 \cdot \sqrt{n\log(n)} \leq \|\mathbf{Q}\boldsymbol{\mu}\|_2^2 = \|\mathbf{Z}\widetilde{\boldsymbol{\mu}}\|_2^2 \leq n\|\widetilde{\boldsymbol{\mu}}\|_2^2 + c_3\|\widetilde{\boldsymbol{\mu}}\|_2^2 \cdot \sqrt{n\log(n)}, \tag{A.8}$$

and

$$\frac{n - c_3\sqrt{n\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda} \cdot \|\widetilde{\boldsymbol{\mu}}\|_2 \leq \boldsymbol{\nu}^\top \mathbf{A}^{-1} \boldsymbol{\nu} \leq \frac{n + c_3\sqrt{n\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} \cdot \|\widetilde{\boldsymbol{\mu}}\|_2$$

Similarly for $\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}$, by Cauchy-Schwarz inequality, for large enough n we have

$$|\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}| = |\mathbf{y}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu}| \leq \|\mathbf{y}\|_2 \cdot \|(\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\mu}\|_2 = \sqrt{n} \cdot \sqrt{\boldsymbol{\mu}^\top \mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-2}\mathbf{Q}\boldsymbol{\mu}}.$$

Applying Lemma A.1 and the inequality (A.8), we have

$$|\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}| \leq \frac{\sqrt{n}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\mathbf{Q}\boldsymbol{\mu}\|_2 \leq \frac{\sqrt{n} \cdot \sqrt{n + c_3\sqrt{n\log(n)}}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\widetilde{\boldsymbol{\mu}}\|_2 \leq \frac{c_4 n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\widetilde{\boldsymbol{\mu}}\|_2,$$

where c_4 is an absolute constant. Note that $\|\widetilde{\boldsymbol{\mu}}\|_2 = \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$. This completes the proof. $\square$

A.2.4 Proof of Lemma A.4

The proof Lemma A.4 utilizes the polarization identity and is similar with the proof of Lemma 4.3.

Proof of Lemma A.4. The proof is very similar to the proof of Lemma 4.3. We have

$$
\begin{aligned}
\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{e}_i y_i &= \frac{1}{4\sqrt{n}}(\mathbf{y}+\sqrt{n}\mathbf{e}_i y_i)^\top \mathbf{A}^{-1}(\mathbf{y}+\sqrt{n}\mathbf{e}_i y_i) - \frac{1}{4\sqrt{n}}(\mathbf{y}-\sqrt{n}\mathbf{e}_i y_i)^\top \mathbf{A}^{-1}(\mathbf{y}-\sqrt{n}\mathbf{e}_i y_i) \\
&\geq \frac{1}{4\sqrt{n}}\bigg[\frac{\|\mathbf{y}+\sqrt{n}\mathbf{e}_i y_i\|_2^2}{\mathrm{tr}(\boldsymbol{\Sigma})+\epsilon_\lambda} - \frac{\|\mathbf{y}-\sqrt{n}\mathbf{e}_i y_i\|_2^2}{\mathrm{tr}(\boldsymbol{\Sigma})-\epsilon_\lambda}\bigg] \\
&= \frac{1}{4\sqrt{n}}\bigg[\frac{2n+2\sqrt{n}}{\mathrm{tr}(\boldsymbol{\Sigma})+\epsilon_\lambda} - \frac{2n-2\sqrt{n}}{\mathrm{tr}(\boldsymbol{\Sigma})-\epsilon_\lambda}\bigg] \\
&= \frac{1}{2\sqrt{n}}\cdot\frac{(n+\sqrt{n})(\mathrm{tr}(\boldsymbol{\Sigma})-\epsilon_\lambda)-(n-\sqrt{n})(\mathrm{tr}(\boldsymbol{\Sigma})+\epsilon_\lambda)}{\mathrm{tr}(\boldsymbol{\Sigma})^2-\epsilon_\lambda^2} \\
&= \frac{1}{2\sqrt{n}}\cdot\frac{2\sqrt{n}\,\mathrm{tr}(\boldsymbol{\Sigma})-2n\epsilon_\lambda}{\mathrm{tr}(\boldsymbol{\Sigma})^2-\epsilon_\lambda^2} \\
&= \frac{\mathrm{tr}(\boldsymbol{\Sigma})-\sqrt{n}\epsilon_\lambda}{\mathrm{tr}(\boldsymbol{\Sigma})^2-\epsilon_\lambda^2},
\end{aligned}
$$

where we use the polarization identity $\mathbf{a}^\top\mathbf{M}\mathbf{b} = 1/4(\mathbf{a}+\mathbf{b})^\top\mathbf{M}(\mathbf{a}+\mathbf{b}) - 1/4(\mathbf{a}-\mathbf{b})^\top\mathbf{M}(\mathbf{a}-\mathbf{b})$ in the first equality and use Lemma A.1 to derive the inequality. This completes the proof. $\square$

B Proof of Theorem 3.2

Here we present the proof of Theorem 3.2.

Proof of Theorem 3.2. By the lower bound of the Gaussian cumulative distribution function (Côté et al., 2012), we have that for any $\boldsymbol{\theta}\in\mathbb{R}^d$,

$$
R(\boldsymbol{\theta}) \geq c_1 \exp\bigg(-\frac{c_2(\boldsymbol{\theta}^\top\boldsymbol{\mu})^2}{\|\boldsymbol{\theta}\|_{\boldsymbol{\Sigma}}^2}\bigg), \tag{B.1}
$$

where $c_1, c_2 > 0$ are absolute constants. By Proposition 4.4, we have

$$
\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}} = \widehat{\boldsymbol{\theta}}_{\mathrm{LS}} = \mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}.
$$

Plugging it into (B.1), we obtain

$$
R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \geq c_1 \exp\bigg\{-\frac{c_2[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2}{\|\mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}^2}\bigg\}. \tag{B.2}
$$

Note that based on our model, we have $\mathbf{X} = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q}$, and

$$
\begin{aligned}
\|\mathbf{X}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}} &= \|(\mathbf{y}\boldsymbol{\mu}^\top+\mathbf{Q})^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}} \\
&\geq \big|\|\boldsymbol{\mu}\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}} - \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}\big| \\
&= \big|\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} - \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}}\big|
\end{aligned} \tag{B.3}
$$

Plugging the above bound into (B.2), we obtain

$$
R(\boldsymbol{\theta}) \geq c_1 \exp\bigg\{-\frac{c_2[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2}{(\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} - \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}})^2}\bigg\}. \tag{B.4}
$$

Denote $\boldsymbol{\nu} = \mathbf{Q}\boldsymbol{\mu}$ and $\mathbf{A} = \mathbf{Q}\mathbf{Q}^\top$. Then by Lemma A.2 and the model definition $\mathbf{X} = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q}$, we have

$$
\begin{aligned}
\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu} &= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}](\mathbf{y}\boldsymbol{\mu}^\top + \mathbf{Q})\boldsymbol{\mu} \\
&= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}](\mathbf{y}\cdot\|\boldsymbol{\mu}\|_2^2 + \mathbf{Q}\boldsymbol{\mu}) \\
&= D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}\mathbf{y}]\cdot\|\boldsymbol{\mu}\|_2^2 \\
&\quad + D^{-1}[(1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu} - \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y}\cdot\boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu}]) \\
&= D^{-1}\cdot[(\|\boldsymbol{\mu}\|_2^2 - \boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} + (1+\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}], \qquad \text{(B.5)}
\end{aligned}
$$

where the third equality follows by the notation $\boldsymbol{\nu} = \mathbf{Q}\boldsymbol{\mu}$. By Lemma A.3 and the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \geq C\max\{\epsilon_\lambda, n\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$ for some large enough constant C, when n is large enough we have

$$
\begin{aligned}
&|\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}| \leq \frac{c_3 n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \leq \frac{2c_4 n}{\mathrm{tr}(\boldsymbol{\Sigma})}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \leq 1, \\
&0 \leq \boldsymbol{\nu}^\top\mathbf{A}^{-1}\boldsymbol{\nu} \leq \frac{n + c_5\sqrt{n\log(n)}}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \leq \frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \leq \frac{2n\|\boldsymbol{\Sigma}\|_2}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_2^2 \leq \frac{1}{2}\cdot\|\boldsymbol{\mu}\|_2^2, \\
&0 \leq \mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} \leq \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} \leq \frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})},
\end{aligned}
$$

where c_3, c_4 are absolute constants. Plugging the bounds above into (A.3), we obtain

$$
\begin{aligned}
|\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}| &\leq D^{-1}\cdot\Big(\|\boldsymbol{\mu}\|_2^2\cdot\mathbf{y}^\top\mathbf{A}^{-1}\mathbf{y} + 2\cdot|\mathbf{y}^\top\mathbf{A}^{-1}\boldsymbol{\nu}|\Big) \\
&\leq D^{-1}\cdot\bigg[\frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_2^2 + \frac{4n}{\mathrm{tr}(\boldsymbol{\Sigma})}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\bigg] \\
&\leq D^{-1}\cdot\frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot(\|\boldsymbol{\mu}\|_2^2 + 2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}) \\
&\leq D^{-1}\cdot\frac{4n}{\mathrm{tr}(\boldsymbol{\Sigma})}\cdot\|\boldsymbol{\mu}\|_2^2,
\end{aligned}
$$

where the last inequality follows by the assumption that $\|\boldsymbol{\mu}\|_2^2 \geq C\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ for some large enough absolute constant C. Therefore we have

$$
[\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\boldsymbol{\mu}]^2 \leq D^{-2}\cdot\frac{n^2}{64[\mathrm{tr}(\boldsymbol{\Sigma})]^2}\cdot\|\boldsymbol{\mu}\|_2^4 = \frac{H(\boldsymbol{\mu},\mathbf{Q},\mathbf{y},\boldsymbol{\Sigma})}{64}\cdot n^2\|\boldsymbol{\mu}\|_2^4, \qquad \text{(B.6)}
$$

where

$$
H(\boldsymbol{\mu},\mathbf{Q},\mathbf{y},\boldsymbol{\Sigma}) := [D\cdot\mathrm{tr}(\boldsymbol{\Sigma})]^{-2} > 0.
$$

We now proceed to study the two terms in the denominator of the exponent in (B.4). We denote

$$
\begin{aligned}
J_1 &= \mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}, \\
J_2 &= \|\mathbf{Q}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}\|_{\boldsymbol{\Sigma}})^2.
\end{aligned}
$$

Then for J_1, with the same derivation as the proof of Lemma 4.7 for I_2, we have

$$J_1 = \sqrt{I_2} \leq 2\sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}.$$

Moreover we also have

$$\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y} = D^{-1} \cdot \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \geq D^{-1} \cdot \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda} \geq (2D)^{-1} \cdot \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma})},$$

where the first inequality follows by Lemma A.3, and the second inequality follows by the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \geq C\epsilon_\lambda$ for some large enough constant C. Then we have

$$J_1 = \mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \geq (2D)^{-1} \cdot \frac{n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}}{\mathrm{tr}(\boldsymbol{\Sigma})} = (1/2) \cdot \sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}},$$

where we use the definition $H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) = [D \cdot \mathrm{tr}(\boldsymbol{\Sigma})]^{-2}$. Therefore in summary we have

$$(1/2) \cdot \sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \leq J_1 \leq 2\sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}, \tag{B.7}$$

where c_5 is an absolute constant. Similarly, for J_2, with the same derivation as the proof of Lemma 4.7 for I_3, we have

$$J_2^2 = I_3 \leq c_5 H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2). \tag{B.8}$$

Moreover, we denote $\mathbf{a} = D^{-1}[(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}) \cdot \mathbf{y} - \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \cdot \boldsymbol{\nu}]$. Then with the same derivation,

$$\begin{aligned} J_2^2 &= \mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{Q}\boldsymbol{\Sigma}\mathbf{Q}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y} \\ &= \mathbf{a}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{Q}\boldsymbol{\Sigma}\mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top)^{-1}\mathbf{a} \\ &= \mathbf{a}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{a}, \end{aligned} \tag{B.9}$$

where we plug in $\boldsymbol{\Sigma} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top$ and $\mathbf{Q} = \mathbf{Z}\boldsymbol{\Lambda}^{1/2}\mathbf{V}^\top$ for $\mathbf{Z}$ with independent sub-Gaussian entries. We have

$$\begin{aligned} J_2^2 &= \mathbf{a}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{Z}\boldsymbol{\Lambda}^2\mathbf{Z}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-1}\mathbf{a} \\ &\geq \mathbf{a}^\top (\mathbf{Z}\boldsymbol{\Lambda}\mathbf{Z}^\top)^{-2}\mathbf{a} \cdot \big[\|\boldsymbol{\Sigma}\|_F^2 - \epsilon_\lambda'\big] \\ &\geq \|\mathbf{a}\|_2^2 \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 - \epsilon_\lambda'}{[\mathrm{tr}(\boldsymbol{\Sigma}) + \epsilon_\lambda]^2} \\ &\geq \|\mathbf{a}\|_2^2 \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 - \epsilon_\lambda'}{2[\mathrm{tr}(\boldsymbol{\Sigma})]^2}. \end{aligned} \tag{B.10}$$

Here the first inequality follows by Lemma A.5, the second inequality follows by Lemma A.1, and the third inequality follows by the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \geq C\epsilon_\lambda$ for some large enough absolute constant C. By the definition of ϵ'_λ in Lemma A.5 and Cauchy-Schwarz inequality, we have

$$\begin{aligned}
\epsilon'_\lambda &:= c_6\big(n \cdot \|\boldsymbol{\Sigma}\|_2^2 + \sqrt{n} \cdot \|\boldsymbol{\Sigma}^2\|_F\big) \\
&\le c_6\big(n \cdot \|\boldsymbol{\Sigma}\|_2^2 + \sqrt{n} \cdot \|\boldsymbol{\Sigma}\|_2 \cdot \|\boldsymbol{\Sigma}\|_F\big) \\
&\le c_6\big(n \cdot \|\boldsymbol{\Sigma}\|_2^2 + 2c_6 n \cdot \|\boldsymbol{\Sigma}\|_2^2 + \|\boldsymbol{\Sigma}\|_F^2/(2c_6)\big) \\
&\le c_7 n \cdot \|\boldsymbol{\Sigma}\|_2^2 + \|\boldsymbol{\Sigma}\|_F^2/2,
\end{aligned}$$

where c_6, c_7 are absolute constants. Plugging the above bound into (B.10) gives

$$J_2^2 \ge \|\mathbf{a}\|_2^2 \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 - c_8 n \cdot \|\boldsymbol{\Sigma}\|_2^2}{4[\mathrm{tr}(\boldsymbol{\Sigma})]^2} \tag{B.11}$$

for some absolute constant c_8. Moreover, by the definition $\mathbf{a}$ and the triangle inequality, we have

$$\begin{aligned}
\|\mathbf{a}\|_2^2 &= \|D^{-1}(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu})\mathbf{y} - \mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} \cdot \boldsymbol{\nu}\|_2^2 \\
&\ge \big[D^{-1}(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu})\|\mathbf{y}\|_2 - D^{-1}(\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y}) \cdot \|\mathbf{Q}\boldsymbol{\mu}\|_2\big]^2 \\
&= D^{-2}\big[(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}) \cdot \sqrt{n} - (\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y}) \cdot \|\mathbf{Q}\boldsymbol{\mu}\|_2\big]^2. \qquad \text{(B.12)}
\end{aligned}$$

By Lemma A.3 and the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \ge C \max\{\epsilon_\lambda, n, n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\}$ for some large enough constant C, we have

$$\begin{aligned}
|\mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}| &\le \frac{c_9 n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \le 1/2, \\
\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y} &\le \frac{n}{\mathrm{tr}(\boldsymbol{\Sigma}) - \epsilon_\lambda} \le \frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})},
\end{aligned}$$

where c_9 is an absolute constant. Moreover, with the same proof as in Lemma A.3, when n is sufficiently large, with probability at least $1 - O(n^{-2})$ we have

$$\|\mathbf{Q}\boldsymbol{\mu}\|_2^2 \le 2n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2.$$

Utilizing these inequalities above, we have

$$\begin{aligned}
&(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}) \cdot \sqrt{n} \ge \sqrt{n}/2, \\
&(\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y}) \cdot \|\mathbf{Q}\boldsymbol{\mu}\|_2 \le \frac{2n}{\mathrm{tr}(\boldsymbol{\Sigma})} \cdot \sqrt{2n}\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \le \sqrt{n}/4,
\end{aligned}$$

where the second line above follows the assumption that $\mathrm{tr}(\boldsymbol{\Sigma}) \ge Cn\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$ for some large enough constant C. Combining these bounds with (B.12), we have

$$\|\mathbf{a}\|_2^2 \ge D^{-2}\big[(1 + \mathbf{y}^\top \mathbf{A}^{-1}\boldsymbol{\nu}) \cdot \sqrt{n} - (\mathbf{y}^\top \mathbf{A}^{-1}\mathbf{y}) \cdot \|\mathbf{Q}\boldsymbol{\mu}\|_2\big]^2 \ge D^{-2}n/16.$$

Further plugging this bound into (B.11), we have

$$J_2^2 \ge \frac{n}{16D^2} \cdot \frac{\|\boldsymbol{\Sigma}\|_F^2 - c_8 n \cdot \|\boldsymbol{\Sigma}\|_2^2}{4[\mathrm{tr}(\boldsymbol{\Sigma})]^2} = H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (c_{10} n \cdot \|\boldsymbol{\Sigma}\|_F^2 - c_{11} n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2), \tag{B.13}$$

where c_{10}, c_{11} are absolute constants, and we use the definition $H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) = [D \cdot \mathrm{tr}(\boldsymbol{\Sigma})]^{-2}$.

Combining (B.8) and (B.13), we obtain

$$H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (c_{10} n \|\boldsymbol{\Sigma}\|_F^2 - c_{11} n^2 \|\boldsymbol{\Sigma}\|_2^2) \leq J_2^2 \leq c_5 H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (n \|\boldsymbol{\Sigma}\|_F^2 + n^2 \|\boldsymbol{\Sigma}\|_2^2). \quad \text{(B.14)}$$

In the rest of the proof, we consider the two cases in Theorem 3.2 separately based on (B.7) and (B.14).

Case 1. Suppose that $n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 \geq C(\|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2)$ for some large enough constant C. Then by (B.7) and (B.14), we have

$$J_1 \geq (1/2) \cdot \sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$$
$$J_2 \leq 2\sqrt{c_5}\sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot \sqrt{n \cdot \|\boldsymbol{\Sigma}\|_F^2 + n^2 \cdot \|\boldsymbol{\Sigma}\|_2^2} \leq (1/4) \cdot \sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}$$

Plugging the above inequalities and (B.6) into (B.4), we obtain Therefore by

$$R(\boldsymbol{\theta}) \geq c_1 \exp\bigg\{ - \frac{c_2 n^2 \|\boldsymbol{\mu}\|_2^4 / 64}{(n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}/4)^2} \bigg\} = c_1 \exp\bigg\{ - \frac{c_{12}\|\boldsymbol{\mu}\|_2^4}{\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2} \bigg\},$$

where c_{12} is an absolute constant. This completes the proof of the first case in Theorem 3.2.

Case 2. Suppose that $\|\boldsymbol{\Sigma}\|_F^2 \geq Cn(\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\boldsymbol{\Sigma}\|_2^2)$ for some large enough constant C. Then by (B.7) we have

$$J_1 \leq 2\sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot n \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} \leq \sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot \sqrt{c_{10}}\|\boldsymbol{\Sigma}\|_F/4. \quad \text{(B.15)}$$

Moreover for J_2, by (B.14) we have

$$J_2^2 \geq H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot (c_{10} n \|\boldsymbol{\Sigma}\|_F^2 - c_{11} n^2 \|\boldsymbol{\Sigma}\|_2^2) \geq H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma}) \cdot c_{10} n \|\boldsymbol{\Sigma}\|_F^2/4,$$

and therefore

$$J_2 \geq \sqrt{H(\boldsymbol{\mu}, \mathbf{Q}, \mathbf{y}, \boldsymbol{\Sigma})} \cdot \sqrt{c_{10} n}\|\boldsymbol{\Sigma}\|_F/2. \quad \text{(B.16)}$$

Plugging (B.6), (B.15) and (B.16) into (B.4), we obtain

$$R(\boldsymbol{\theta}) \geq c_1 \exp\bigg\{ - \frac{c_2 n^2 \|\boldsymbol{\mu}\|_2^4 / 64}{(\sqrt{c_{10} n}\|\boldsymbol{\Sigma}\|_F/4)^2} \bigg\} = c_1 \exp\bigg\{ - \frac{c_{13} n\|\boldsymbol{\mu}\|_2^4}{\|\boldsymbol{\Sigma}\|_F^2} \bigg\},$$

where c_{13} is an absolute constant. This completes the proof of the second case in Theorem 3.2. ☐

C Proof of Corollaries

Here we provide the proof of the Corollaries 3.3, 3.6 and 3.8 in Section 3.

C.1 Proof of Corollary 3.3

The proof of Corollary 3.3 is a direct application of Theorem 3.1. The detailed proof is as follows.

Proof of Corollary 3.3. When $\boldsymbol{\Sigma} = \mathbf{I}$, we have $\mathrm{tr}(\boldsymbol{\Sigma}) = d$, $\|\boldsymbol{\Sigma}\|_2 = 1$, $\|\boldsymbol{\Sigma}\|_F = \sqrt{d}$ and $\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}} = \|\boldsymbol{\mu}\|_2$. Under the condition in Corollary 3.3 that $d \geq C \max\big\{n^2, n\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_2\big\}$ and $\|\boldsymbol{\mu}\|_2 \geq C$ for some large enough absolute constant C, it is easy to check that the conditions of Theorem 3.1

$$\mathrm{tr}(\boldsymbol{\Sigma}) = \Omega\big(\max\big\{n^{3/2}\|\boldsymbol{\Sigma}\|_2, n\|\boldsymbol{\Sigma}\|_F, n\sqrt{\log(n)}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}\big\}\big),\ \|\boldsymbol{\mu}\|_2 \geq C\|\boldsymbol{\Sigma}\|_2$$

hold. Therefore by Theorem 3.1, we have

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \leq \exp\left(\frac{-c_1 n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \|\boldsymbol{\Sigma}\|_F^2 + n\|\boldsymbol{\Sigma}\|_2^2}\right) \leq \exp\left(\frac{-c_2 n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_2^2 + d}\right),$$

where c_1, c_2 are absolute constants. This completes the proof. □

C.2 Proof of Corollary 3.6

Here we present the proof of Corollary 3.6, which is mostly based on the estimation of the order of the summations $\sum_{k=1}^d k^{-\alpha}$ and $\sum_{k=1}^d k^{-2\alpha}$. We first present the full version of the corollary with detailed dependency in the sample size n as follows.

Corollary C.1. [Full version of Corollary 3.6] Suppose that $\lambda_k = k^{-\alpha}$, and one of the following conditions hold:

1. $\alpha \in [0, 1/2)$, $d = \widetilde{\Omega}(n^{\frac{3}{2(1-\alpha)}} + n^2 + (n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$, and $\|\boldsymbol{\mu}\|_2 = \omega(1 + n^{-1/4}d^{1/4-\alpha/2})$.
2. $\alpha = 1/2$, $d = \widetilde{\Omega}(n^3 + n^2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2)$, and $\|\boldsymbol{\mu}\|_2 = \omega(1 + n^{-1/4}(\log(d))^{1/4})$.
3. $\alpha \in (1/2, 1)$, $d = \widetilde{\Omega}(n^{\frac{3}{2(1-\alpha)}} + (n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$, and $\|\boldsymbol{\mu}\|_2 = \omega(1)$.

Then with probability at least $1-n^{-1}$, the population risk of the maximum margin classifier satisfies $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) = o(1)$.

Proof of Corollary C.1. We first consider the case when $\alpha \in [0, 1/2)$. We have

$$\mathrm{tr}(\boldsymbol{\Sigma}) = \sum_{k=1}^d \lambda_k = \sum_{k=1}^d k^{-\alpha} \geq \int_{t=1}^d t^{-\alpha}\mathrm{d}t = \frac{d^{1-\alpha}}{1-\alpha} - \frac{1}{1-\alpha} > \frac{d^{1-\alpha}}{2(1-\alpha)}$$

when d is sufficiently large. Similarly, we have

$$\|\boldsymbol{\Sigma}\|_F^2 = \sum_{k=1}^d \lambda_k^2 = 1 + \sum_{k=2}^d k^{-2\alpha} \leq 1 + \int_{t=1}^{d-1} t^{-2\alpha}\mathrm{d}t = 1 + \frac{(d-1)^{1-2\alpha}}{1-2\alpha} - \frac{1}{1-2\alpha} \leq 1 + \frac{d^{1-2\alpha}}{1-2\alpha}.$$

Therefore, a sufficient condition for the assumptions in Theorem 3.1 to hold is that $\|\boldsymbol{\mu}\|_2 = \omega(1)$ and

$$\frac{d^{1-\alpha}}{2(1-\alpha)} \geq Cn^{3/2},$$
$$\frac{d^{1-\alpha}}{2(1-\alpha)} \geq Cn\cdot\sqrt{1 + \frac{d^{1-2\alpha}}{1-2\alpha}},$$
$$\frac{d^{1-\alpha}}{2(1-\alpha)} \geq Cn\sqrt{\log(n)}\cdot\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}.$$

After simplifying the result, we derive the condition that $d = \widetilde{\Omega}(n^{\frac{3}{2(1-\alpha)}} + n^2 + (n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$. We further check the conditions on $\|\boldsymbol{\mu}\|_2$ that lead to $o(1)$ population risk. Note that when $\|\boldsymbol{\mu}\|_2 = \omega(1)$, $\|\boldsymbol{\mu}\|_2^4/\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 = \omega(1)$. We also check the condition that $n\|\boldsymbol{\mu}\|_2^4/\|\boldsymbol{\Sigma}\|_F^2 = \omega(1)$. A sufficient condition is that

$$n\|\boldsymbol{\mu}\|_2^4 = \omega\bigg(1 + \frac{d^{1-2\alpha}}{1-2\alpha}\bigg).$$

Simplifying the condition completes the proof for the case $\alpha \in [0, 1/2)$.

For the case $\alpha = 1/2$, we have

$$\mathrm{tr}(\boldsymbol{\Sigma}) = \sum_{k=1}^{d} \lambda_k = \sum_{k=1}^{d} k^{-1/2} \geq \int_{t=1}^{d} t^{-1/2} \mathrm{d}t = \frac{d^{1-1/2}}{1-1/2} - \frac{1}{1-1/2} > \sqrt{d}$$

when d is sufficiently large. Moreover,

$$\|\boldsymbol{\Sigma}\|_F^2 = \sum_{k=1}^{d} \lambda_k^2 = 1 + \sum_{k=2}^{d} k^{-1} \leq 1 + \int_{t=1}^{d-1} t^{-1} \mathrm{d}t = 1 + \log(d-1) \leq 1 + \log(d).$$

Verifying the conditions

$$\begin{aligned}
&\sqrt{d} \geq Cn^{3/2}, \\
&\sqrt{d} \geq Cn \cdot \sqrt{1 + \log(d)}, \\
&\sqrt{d} \geq Cn\sqrt{\log(n)} \cdot \|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}
\end{aligned}$$

then gives a sufficient condition $d = \widetilde{\Omega}(n^3 + n^2\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2)$, $\|\boldsymbol{\mu}\|_2 = \omega(1)$ for the assumptions in Theorem 3.1 to hold. It is also easy to verify that when $\|\boldsymbol{\mu}\|_2 = \omega(1 + n^{-1/4}(\log(d))^{1/4})$ we have $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) = o(1)$.

Finally for the case $\alpha \in (1/2, 1)$, we have

$$\mathrm{tr}(\boldsymbol{\Sigma}) = \sum_{k=1}^{d} \lambda_k = \sum_{k=1}^{d} k^{-\alpha} \geq \int_{t=1}^{d} t^{-\alpha} \mathrm{d}t = \frac{d^{1-\alpha}}{1-\alpha} - \frac{1}{1-\alpha}.$$

Moreover, in this setting we have $\|\boldsymbol{\Sigma}\|_F^2 \leq c_1$ for some absolute constant c_1. It is therefore easy to check that $\|\boldsymbol{\mu}\|_2 = \omega(1)$ and

$$d = \widetilde{\Omega}(n^{\frac{3}{2(1-\alpha)}} + (n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}})^{\frac{1}{1-\alpha}})$$

are sufficient for the assumptions in Theorem 3.1 to hold, and we also have $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) = o(1)$. $\square$

C.3 Proof of Corollary 3.8

The proof of Corollary 3.8 for the rare/weak feature model is rather straightforward.

Proof of Corollary 3.8. Note that in the rare/weak feature model we have $\|\boldsymbol{\mu}\|_2 = \gamma\sqrt{s}$. Therefore the conditions of Corollary 3.3 are satisfied and we have

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \le \exp\bigg(-\frac{c_1 n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_2^2 + d}\bigg) = \exp\bigg(-\frac{c_1 n\gamma^4 s^2}{n\gamma^2 s + d}\bigg),$$

where c_1 is an absolute constant. This completes the proof. $\square$

D Experiments

In this section we present simulation results to backup our population risk bound in Theorem 3.1. We generate $\mathbf{u}$ as a standard Gaussian vector, and set $\boldsymbol{\Sigma} = \mathrm{diag}\{\lambda_1, \dots, \lambda_d\}$ with $\lambda_k = k^{-\alpha}$ for some parameter $\alpha \in [0, 1)$, which matches the setting studied in Section 3. The mean vector $\boldsymbol{\mu}$ is generated uniformly from the sphere centered at the origin with radius r. All population risks are calculated by taking the average of 100 independent experiments. Note that since we are considering Gaussian mixtures in our experiments, the population risk can be directly calculated with the Gaussian cumulative distribution function:

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) = \mathbb{P}[\boldsymbol{\theta}^\top \boldsymbol{\mu} < y \cdot \widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}^\top \boldsymbol{\Lambda}^{1/2} \mathbf{u}].$$

The derivation of the above result is in the proof of Lemma 4.5 in Appendix A.1.3.

Population risk versus the norm of the mean vector $\|\boldsymbol{\mu}\|_2$. We first present experimental results on the relation between the population risk and the norm of the mean vector $\|\boldsymbol{\mu}\|_2$. Note that in our setting, the risk bound in Theorem 3.1 reduces to the following bound:

$$R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}) \le \exp\Bigg(\frac{-C' n\|\boldsymbol{\mu}\|_2^4}{n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2 + \sum_{k=1}^d k^{-2\alpha}}\Bigg).$$

Based on this bound, we can first see that the population risk should be smaller when α is larger. Moreover, the dependency of $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ depends on the comparison between the scaling of the two terms in the denominator. When

$$\sum_{k=1}^d k^{-2\alpha} \ge n\|\boldsymbol{\mu}\|_{\boldsymbol{\Sigma}}^2, \tag{D.1}$$

we can expect that $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ should be roughly of order $\|\boldsymbol{\mu}\|_2^4$. On the other hand, if (D.1) does not hold, then $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ should be roughly of order $\|\boldsymbol{\mu}\|_2^2$. It is also clear that whether (D.1) holds heavily depends on the values of the sample size n and α: when n is large, then (D.1) is less likely to be satisfied. Moreover, when $\alpha > 1/2$, (D.1) cannot hold because in this case $\sum_{k=1}^d k^{-2\alpha}$ is upper bounded by a constant.

In Figure 2, we verify the above argument by verifying the dependency of the population risk $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ on the norm of the mean vector $\|\boldsymbol{\mu}\|_2$ with different values of α and sample size n. From Figures 2(a) and 2(c), we can see that $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ decreases with $\|\boldsymbol{\mu}\|_2$ and α. From 2(b), we verify that when $n = 10$ (which is rather small) and when $\alpha = 0, 0.2, 0.4$, $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ is linear in $\|\boldsymbol{\mu}\|_2^2$. This verifies our discussion for the setting when (D.1) holds. On the other hand, when $\alpha = 0.6, 0.8$, $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ has a higher order dependency in $\|\boldsymbol{\mu}\|_2^2$, which is because $\sum_{k=1}^d k^{-2\alpha}$ is upper bounded by a constant and (D.1) cannot hold. In Figure 2(d), we further verify that when

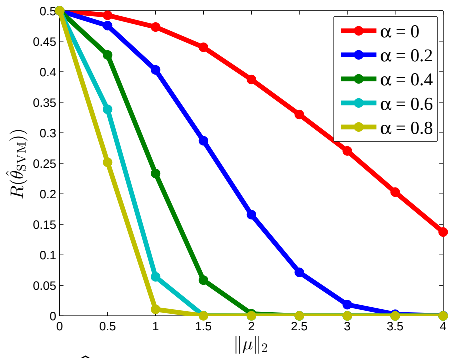

(a) $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ versus $\|\boldsymbol{\mu}\|_2$, $n = 10$, $d = 2000$

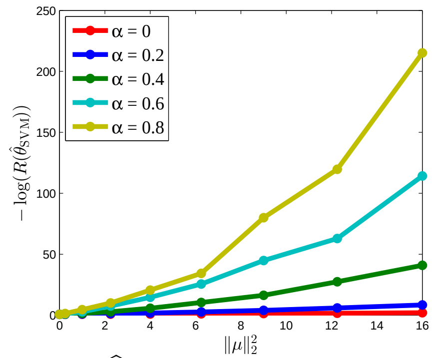

(b) $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ versus $\|\boldsymbol{\mu}\|_2^2$, $n = 10$, $d = 2000$

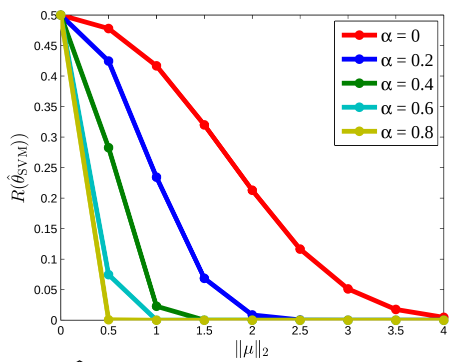

(c) $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ versus $\|\boldsymbol{\mu}\|_2$, $n = 100$, $d = 2000$

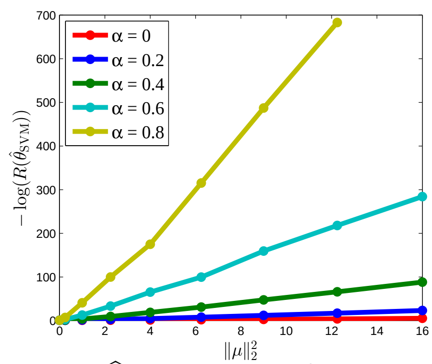

(d) $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ versus $\|\boldsymbol{\mu}\|_2^2$, $n = 100$, $d = 2000$

Figure 2: Experiments on the dependency of the population risk $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ on the norm of the mean vector $\|\boldsymbol{\mu}\|_2$ with different values of α and sample size n. (a) and (b) gives the curves with $n = 10$, while (c) and (d) are for the case $n = 100$. Moreover, (a) and (c) gives the curves of $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ versus $\|\boldsymbol{\mu}\|_2$, and to further test the tightness of our risk bound, in (c) and (d) we also study the relation between $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ and $\|\boldsymbol{\mu}\|_2^2$. The dimension d is set to 2000 in all these figures. In (d) we omit the last point $\|\boldsymbol{\mu}\|_2 = 16$ in the curve for $\alpha = 0.8$ because the population risk in this case is too small and is dominated by the numerical accuracy.

$n = 100$, (D.1) never hold and $-\log(R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}}))$ is of order $\|\boldsymbol{\mu}\|_2^2$ for all choices of α. This set of experiments verifies our risk bound in Theorem 3.1.

Verification of the dimension-dependent and dimension-free settings. In Corollary 3.6, we have discussed that when $\alpha < 1/2$, achieving a small population risk requires a larger $\|\boldsymbol{\mu}\|_2$ when d is larger. On the other hand, when $\alpha > 1/2$, the requirement on $\|\boldsymbol{\mu}\|_2$ to achieve small population error is dimension-free. Here we present experimental results to verify our claim. The results are given in Figure 3. We can see very clearly that when $\alpha = 0.2$, the risk curves for different d are different, and larger d results in worse population risk. However, when $\alpha = 0.8$, all the risk curves are almost exactly the same, which indicates that the population risk is dimension-free. This verifies our claim in Corollary 3.6.

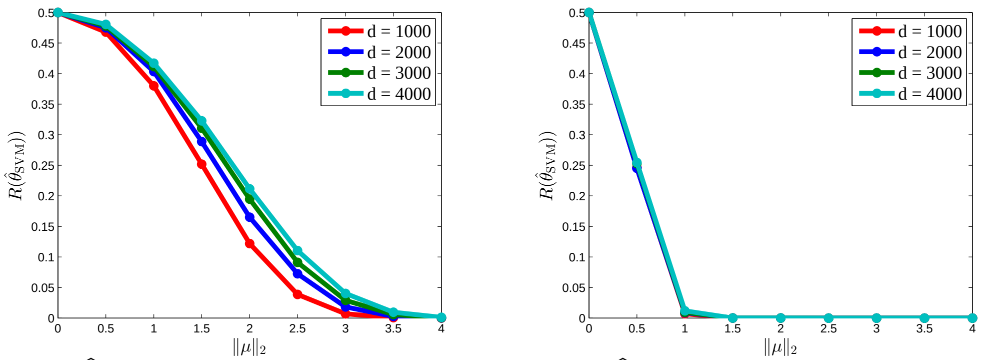

(a) $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ versus $\|\boldsymbol{\mu}\|_2$, $\alpha = 0.2$, $n = 10$ (b) $R(\widehat{\boldsymbol{\theta}}_{\mathrm{SVM}})$ versus $\|\boldsymbol{\mu}\|_2$, $\alpha = 0.8$, $n = 10$

Figure 3: The population risk curve with respect to $\|\boldsymbol{\mu}\|_2$ with different values of α and dimension d. (a) shows the result for $\alpha = 0.2$, while (b) is for the case $\alpha = 0.8$. The sample size n is set to 10 in both experiments.

References

Arora, S., Cohen, N., Hu, W. and Luo, Y. (2019). Implicit regularization in deep matrix factorization. *Advances in Neural Information Processing Systems* **32**.

Bartlett, P. L., Long, P. M., Lugosi, G. and Tsigler, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences* .

Belkin, M., Hsu, D., Ma, S. and Mandal, S. (2019a). Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences* **116** 15849–15854.

Belkin, M., Hsu, D. and Xu, J. (2019b). Two models of double descent for weak features. *arXiv preprint arXiv:1903.07571* .

Belkin, M., Ma, S. and Mandal, S. (2018). To understand deep learning we need to understand kernel learning. In *International Conference on Machine Learning*. PMLR.

Chatterji, N. S. and Long, P. M. (2020). Finite-sample analysis of interpolating linear classifiers in the overparameterized regime. *arXiv preprint arXiv:2004.12019* .

Côté, F. D., Psaromiligkos, I. N. and Gross, W. J. (2012). A chernoff-type lower bound for the gaussian q-function. *arXiv preprint arXiv:1202.6483* .

Donoho, D. and Jin, J. (2008). Higher criticism thresholding: Optimal feature selection when useful features are rare and weak. *Proceedings of the National Academy of Sciences* **105** 14790–14795.

Gunasekar, S., Lee, J., Soudry, D. and Srebro, N. (2018). Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*.

Gunasekar, S., Woodworth, B. E., Bhojanapalli, S., Neyshabur, B. and Srebro, N. (2017). Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*.

Hastie, T., Montanari, A., Rosset, S. and Tibshirani, R. J. (2019). Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560* .

Hsu, D., Muthukumar, V. and Xu, J. (2020). On the proliferation of support vectors in high dimensions. *arXiv preprint arXiv:2009.10670* .

Ji, Z. and Telgarsky, M. (2019). The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory*.

Jin, J. (2009). Impossibility of successful classification when useful features are rare and weak. *Proceedings of the National Academy of Sciences* **106** 8859–8864.

Liao, Z., Couillet, R. and Mahoney, M. (2020). A random matrix analysis of random fourier features: beyond the gaussian kernel, a precise phase transition, and the corresponding double descent. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*.

Lyu, K. and Li, J. (2019). Gradient descent maximizes the margin of homogeneous neural networks. *arXiv preprint arXiv:1906.05890* .

Mei, S. and Montanari, A. (2019). The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355* .

Montanari, A. and Zhong, Y. (2020). The interpolation phase transition in neural networks: Memorization and generalization under lazy training. *arXiv preprint arXiv:2007.12826* .

Muthukumar, V., Narang, A., Subramanian, V., Belkin, M., Hsu, D. and Sahai, A. (2020a). Classification vs regression in overparameterized regimes: Does the loss function matter? *arXiv preprint arXiv:2005.08054* .

Muthukumar, V., Vodrahalli, K., Subramanian, V. and Sahai, A. (2020b). Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory* **1** 67–83.

Nacson, M. S., Srebro, N. and Soudry, D. (2019). Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate. In *The 22nd International Conference on Artificial Intelligence and Statistics*.

Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S. and Srebro, N. (2018). The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research* **19** 2822–2878.

Tsigler, A. and Bartlett, P. L. (2020). Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286* .

Vershynin, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027* .

Wang, K. and Thrampoulidis, C. (2020). Benign overfitting in binary classification of gaussian mixtures. *arXiv preprint arXiv:2011.09148* .

Wu, D. and Xu, J. (2020). On the optimal weighted ℓ_2 regularization in overparameterized linear regression. *Advances in Neural Information Processing Systems* **33**.