

Bayesian Sparse Gaussian Mixture Model for Clustering in High Dimensions

Dapeng Yao

DYAO10@JHU.EDU

*Department of Applied Mathematics and Statistics, Johns Hopkins University,
Baltimore, Maryland, U.S.A.*

Fangzheng Xie

FXIE@IU.EDU

*Department of Statistics, Indiana University,
Bloomington, Indiana, U.S.A.*

Yanxun Xu

YANXUN.XU@JHU.EDU

*Department of Applied Mathematics and Statistics, Johns Hopkins University,
Baltimore, Maryland, U.S.A.*

Editor: Debdeep Pati

Abstract

We study the sparse high-dimensional Gaussian mixture model when the number of clusters is allowed to grow with the sample size. A minimax lower bound for parameter estimation is established, and we show that a constrained maximum likelihood estimator achieves the minimax lower bound. However, this optimization-based estimator is computationally intractable because the objective function is highly nonconvex and the feasible set involves discrete structures. To address the computational challenge, we propose a computationally tractable Bayesian approach to estimate high-dimensional Gaussian mixtures whose cluster centers exhibit sparsity using a continuous spike-and-slab prior. We further prove that the posterior contraction rate of the proposed Bayesian method is minimax optimal. The misclustering rate is obtained as a by-product using tools from matrix perturbation theory. The proposed Bayesian sparse Gaussian mixture model does not require pre-specifying the number of clusters, which can be adaptively estimated. The validity and usefulness of the proposed method is demonstrated through simulation studies and the analysis of a real-world single-cell RNA sequencing data set.

Keywords: Clustering, High dimensions, Minimax estimation, Posterior contraction, Single-cell sequencing.

1. Introduction

Clustering is a powerful tool for detecting structures in heterogeneous data and identifying homogeneous subgroups with a wide range of applications, such as genomics (Gu and Liu, 2008), pattern recognition (Diday et al., 1981), and topic modeling (Blei et al., 2003). In many scientific domains, data are often high-dimensional, i.e., the dimension of observations can be larger than the sample size. For example, an important task in the analysis of single-cell RNA-sequencing data, where the number of genes (dimension) is usually larger than the number of cells (sample size), is to cluster cells and identify functional cell subpopulations (Cao et al., 2019). A principal challenge of extending the low-dimensional clustering techniques to high dimensions is the well-known “curse of dimensionality.” To overcome

this issue, dimensionality reduction (Ding et al., 2002) or additional structural assumptions (Cai et al., 2019) are usually necessary in high dimensional clustering methods.

High dimensional clustering and mixture models have attracted attention recently from the frequentist perspective. When the dimension p has at most the same order as the sample size n , Doss et al. (2023) studied the optimal rate of estimation in a finite Gaussian location mixture model without a separation condition. Löffler et al. (2021) showed that spectral clustering is minimax optimal in the Gaussian mixture model with isotropic covariance matrix when $p = o(n\Delta)$, where Δ is the minimal distance among cluster centers. When $p \gg n$, Azizyan et al. (2013) considered a simple case in which there are only two clusters with equal mixing weights and same isotropic covariance matrices. Jin and Wang (2016) and Jin et al. (2017) proposed influential features principal component analysis based on feature selection and principal component analysis. A phase transition phenomenon in high dimensional clustering problem was also investigated in Jin and Wang (2016) and Jin et al. (2017) across different sparsity and signal levels. Cai et al. (2019) proposed a modified Expectation-Maximization (EM) algorithm based on sparse discriminant vectors to obtain the minimax optimal convergence rate of the excess mis-clustering error. In terms of density estimation, Ashtiani et al. (2020) obtained a near-optimal convergence rate for high dimensional location-scale mixtures with respect to the total variation distance.

Despite these theoretical and computational developments in high dimensional clustering, most frequentist approaches dealing with finite mixtures assume that the number of clusters K is either known or needs to be estimated consistently using techniques such as cross-validation (Smyth, 2000) and the gap statistics (Tibshirani et al., 2001). In contrast, Bayesian methods treat K as an unknown parameter and put a prior on it. For example, Miller and Harrison (2018) proposed a mixture of finite mixtures model with a Gibbs sampler, and the posterior consistency of K was later studied in Miller (2023). Ohn and Lin (2022) established a near optimal rate for estimating finite Gaussian mixtures with respect to the Wasserstein distance when K is unknown and allowed to grow with n . In the context of Bayesian model-based clustering for high-dimensional data, Tadesse et al. (2005) proposed a method for uncovering cluster structure and identifying the discriminating variables simultaneously and exploited reversible jump Markov chain Monte Carlo (MCMC) for posterior sampling. Gao et al. (2020) proposed a Bayesian structured linear model, which encompasses the bi-clustering problem as a special case, using a subset-selection prior and establishing an optimal posterior contraction rate for parameter estimation. Chandra et al. (2022) proposed a Bayesian latent factor mixture model and investigated the behavior of the induced cluster memberships as p goes to infinity whereas n remains fixed. However, a general theoretical framework for Bayesian analysis of high-dimensional clustering in terms of both parameter estimation and mis-clustering error is yet to be established.

The Gaussian mixture model we consider lies in the regime of high dimensionality with sparsity structures. There has been a growing interest in Bayesian inference with sparsity-enforcing priors. One commonly-used prior is the spike-and-slab prior (Mitchell and Beauchamp, 1988), which is a mixture of a point mass at zero and a relatively “flat” absolutely continuous density. The spike-and-slab LASSO prior (Ročková and George, 2018) borrows the similarity between the LASSO and Laplace prior, and combines it with a continuous version of the spike-and-slab prior. Theoretical properties of the spike-and-slab LASSO prior were well studied in the context of regression models, graphical models,

and Gaussian sequence models (see Bai et al., 2021 for a review). Another class of sparsity-enforcing priors is global-local shrinkage priors, such as the horseshoe prior (Carvalho et al., 2009) and the Dirichlet-Laplace prior (Bhattacharya et al., 2015). We refer the readers to Tadesse and Vannucci (2021) and references therein. However, when these types of priors are adapted to high-dimensional sparse clustering problems with an unknown number of clusters, their theoretical properties remain unclear.

This paper presents the Bayesian analysis of a high-dimensional sparse Gaussian mixture model using a spike-and-slab LASSO prior and establishes the optimality of the proposed estimation procedure. Our main contribution is threefold. First, we fully characterize the minimax rate for parameter estimation in the high-dimensional sparse Gaussian mixture model, in which the number of clusters is allowed to grow with sample size. A frequentist constrained maximum likelihood estimator (MLE) is proved to achieve minimax optimality. Second, since the frequentist optimization-based estimator is computationally intractable, we propose a Bayesian sparse Gaussian mixture model to address this computational challenge, and establish its theoretical properties by showing that the posterior contraction rate for parameter estimation is minimax optimal. Third, we obtain a posterior contraction rate for the mis-clustering error using tools from matrix perturbation theory (Yu et al., 2014). To the best of our knowledge, this paper represents the first effort in establishing the theoretical results of parameter estimation and clustering recovery in high-dimensional Bayesian sparse Gaussian mixture model with a growing number of clusters.

The rest of this paper is organized as follows. In Section 2, we introduce the high-dimensional clustering problem and our model, establish the minimax lower bound for parameter estimation, and propose a frequentist constrained MLE that achieves the minimax lower bound. Section 3 elaborates on the main theoretical results, including the optimal posterior contraction rate and the mis-clustering error. We demonstrate the practical performance of the proposed method through simulation studies in Section 4 and a real-world application to clustering single-cell RNA sequencing data in Section 5.

Notations: Let $|S|$ denote the cardinality of S if the set S is finite or the volume (Lebesgue measure) of S if S is a Lebesgue-measurable infinite subset in Euclidean space. Denote $[n]$ as the set of all consecutive integers $\{1, \dots, n\}$. We use \lesssim and \gtrsim to denote the inequality up to a constant. In other words, $a \lesssim$ (\gtrsim , resp.) b if $a \leq$ (\geq , resp.) Cb for some constant $C > 0$. We write $a \asymp b$ if $a \lesssim b$ and $b \lesssim a$. We use $\lfloor x \rfloor$ to denote the greatest integer less than or equal to the real number x and $\lceil x \rceil$ to denote the smallest integer greater than or equal to the real number x . For a p -dimensional vector $\mathbf{x} \in \mathbb{R}^p$, we denote x_i as the i th coordinate of x . Also, we denote $\|\mathbf{x}\|_1 = \sum_{i=1}^p |x_i|$ as the ℓ_1 -norm, $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^p |x_i|^2}$ as the ℓ_2 -norm, $\|\mathbf{x}\|_\infty = \max_{i \in [p]} |x_i|$ as the ℓ_∞ -norm, and $\|\mathbf{x}\|_0 = \sum_{i=1}^p \mathbb{1}(x_i \neq 0)$. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, let A_{ij} denote the (i, j) -entry of \mathbf{A} and let \mathbf{A}_i and \mathbf{A}_j be the i th row and j th column of \mathbf{A} , respectively. We denote $\|\mathbf{A}\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^m |A_{ij}|^2}$ to be the Frobenius norm of \mathbf{A} and $\|\mathbf{A}\|_2$ to be the spectral norm of \mathbf{A} . We denote $\mathbf{A} \succ \mathbf{0}$ if \mathbf{A} is a positive definite matrix. The prior and posterior probability distributions are denoted as Π and the corresponding densities with respect to the underlying σ -finite measure (whenever it exists) are denoted as π . We denote $D_{KL}(\mathbb{P} \|\mathbb{Q})$ the Kullback–Leibler divergence between any probability measures \mathbb{P} and \mathbb{Q} . The ϵ -packing number of a metric space Θ with respect to the metric d , which is the maximum number of pairwise disjoint balls contained in Θ

with radii ϵ , is denoted as $M(\epsilon, \Theta, d)$. In the rest of the paper, we will use an asterisk to represent the ground true values of the parameters that give rise to the data distribution.

2. Model

This section introduces a Gaussian mixture model for high-dimensional clustering, establishing a minimax lower bound and proposing a Bayesian approach for tractable inference.

2.1 Gaussian mixture model and clustering

Let $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_n]$ be a $p \times n$ data matrix, where rows represent variables or features, and columns represent observations. We assume that the data exhibits a clustering structure that can be described through a Gaussian mixture model as follows. Let $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K \in \mathbb{R}^p$ be the cluster centroids of the respective clusters, where $K \geq 1$ is the number of clusters. Let $\mathbf{z} = (z_1, \dots, z_n)^T \in [K]^n$ be the cluster membership vector for observations, with $z_i = k$ indicating that \mathbf{Y}_i belongs to the k th cluster. The distribution of \mathbf{Y}_i is given by

$$\mathbf{Y}_i = \boldsymbol{\mu}_{z_i} + \boldsymbol{\epsilon}_i, \quad (1)$$

where $\boldsymbol{\epsilon}_i \sim N_p(\mathbf{0}, \boldsymbol{\Sigma})$ independently for $i \in [n]$. The goal is to estimate the cluster centroids $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K$ as well as to recover the latent cluster membership vector \mathbf{z} .

This paper considers the asymptotic regime where both p and n go to infinity and $p/n \rightarrow \infty$. When p does not exceed n , Azizyan et al. (2013) proved that the expected clustering accuracy (which will be defined formally later) depends on the dimension p through the rate $\sqrt{p/n}$ in the two-cluster problem without additional structural assumptions. Under the regime that $p/n \rightarrow \infty$ considered in our framework, we posit the following sparse structure on the cluster mean vectors $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K$. Denote $\boldsymbol{\mu} = [\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K]$ as the matrix concatenated by the mean vectors of all clusters and define the support of $\boldsymbol{\mu}$ as $\text{supp}(\boldsymbol{\mu}) = \{j \in [p] : (\boldsymbol{\mu}_j)^T \neq \mathbf{0}\}$. We say that $\boldsymbol{\mu}$ is jointly s -sparse if $|\text{supp}(\boldsymbol{\mu})| \leq s$. Moreover, we require that not only each $\boldsymbol{\mu}_k$ has at most s_n non-zero coordinates, namely, $\sum_{j=1}^p \mathbb{1}(\mu_{kj} \neq 0) \leq s_n$ for all $k \in [K]$, but also that $\boldsymbol{\mu}$ is jointly s_n -sparse. We assume $s_n \rightarrow \infty$ as $n \rightarrow \infty$. In the sequel, we will drop the subscript n from s_n and write $s = s_n$ for notation simplicity, but the readers should be reminded that s depends on n implicitly.

Denote $\mathbf{e}_i \in \mathbb{R}^K$ the unit vector that has value 1 at the i th coordinate and 0 elsewhere. Let $\mathbf{L} = [\mathbf{l}_1, \dots, \mathbf{l}_n]^T \in \mathbb{R}^{n \times K}$ where $\mathbf{l}_i = \mathbf{e}_{z_i}$. Then \mathbf{L} is the matrix whose rows represent cluster memberships of the n observations. It follows immediately that the expected data matrix can be written as $\mathbb{E}(\mathbf{Y}) = \boldsymbol{\mu}\mathbf{L}^T$. Namely, our model can be represented as a signal-plus-noise model matrix $\mathbf{Y} = \boldsymbol{\mu}\mathbf{L}^T + \mathbf{E}$, where $\mathbf{E} = [\boldsymbol{\epsilon}_1, \dots, \boldsymbol{\epsilon}_n]$ is the mean-zero noise matrix where $\boldsymbol{\epsilon}_i$'s are independent normal random vectors with mean zero and covariance matrix $\boldsymbol{\Sigma}$. As K is typically much smaller than n , the above representation of the model is similar to those in Cape et al. (2018) and Agterberg et al. (2022) because the data matrix has a low expected rank. Nevertheless, the sparse Gaussian mixture model differs from Cape et al. (2018) and Agterberg et al. (2022) in that the columns of the expected data matrix have the clustering structure and the rows have the sparsity structure. Following the previous convention of using asterisk to indicate true parameter values, we denote $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)$ the underlying truth of $(\boldsymbol{\mu}, \mathbf{L}, \boldsymbol{\Sigma})$ throughout the rest of the paper.

2.2 Minimax lower bound

One of the major theoretical contributions of this paper is to study the estimation error of the mean matrix $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$ in the proposed sparse Gaussian mixture model. This differs from most existing minimax results in the clustering literature, which predominantly focus on evaluating the mis-clustering error. As the first step towards the complete theory, we establish the minimax lower bound. Formally, consider the following parameter space

$$\Theta_K^* = \{(\boldsymbol{\mu}, \mathbf{L}, \boldsymbol{\Sigma}) : \boldsymbol{\mu} \in \mathbb{R}^{p \times K}, \mathbf{L} \in \mathcal{L}_K, |\text{supp}(\boldsymbol{\mu})| \leq s, \|\boldsymbol{\mu}\mathbf{L}^T\|_F^2 = O(sn), \\ \boldsymbol{\Sigma} \succ \mathbf{0}, 0 < m_{\boldsymbol{\Sigma}} \leq \lambda_{\min}(\boldsymbol{\Sigma}) \leq \lambda_{\max}(\boldsymbol{\Sigma}) \leq M_{\boldsymbol{\Sigma}} < \infty\},$$

where $\mathcal{L}_K = \{\mathbf{L} \in \mathbb{R}^{n \times K} : \mathbf{L} = [\mathbf{l}_1, \dots, \mathbf{l}_n]^T, \mathbf{l}_i \in \{0, 1\}^K, \|\mathbf{l}_i\|_0 = 1 \text{ for all } i \in [n]\}$ is the set of cluster assignment matrices, $\text{supp}(\boldsymbol{\mu})$ is the set of indices of the non-zero rows of $\boldsymbol{\mu}$, and $\lambda_{\min}(\boldsymbol{\Sigma})$ and $\lambda_{\max}(\boldsymbol{\Sigma})$ represents the smallest and largest singular value of $\boldsymbol{\Sigma}$ respectively. We also denote $\Delta = \min_{k_1 \neq k_2} \|\boldsymbol{\mu}_{k_1}^* - \boldsymbol{\mu}_{k_2}^*\|_2$ as the minimum separation of the cluster centers.

We next present a collection of assumptions that are necessary in theoretical analyses.

Assumption 1 (*Low rank*) $K \log n \lesssim \log p$, $K \leq s$.

Assumption 2 (*Minimum separation*) $\Delta \geq 1/n^q$ for some constant $q > 0$.

Assumption 3 (*High dimensionality*) $p/n \rightarrow \infty$.

Assumption 1 is a mild low-rank assumption and can be satisfied even with increasing K . Assumption 2 requires that the centers of different clusters are well separated and is common in high-dimensional clustering problems. It also guarantees the identifiability of $(\boldsymbol{\mu}, \mathbf{L})$ up to a permutation. Assumption 3 requires $p/n \rightarrow \infty$ and it describes the high-dimensional nature of the problem. Below, Theorem 1 establishes the minimax lower bound for estimating the mean matrix with regard to the Frobenius norm.

Theorem 1 *Let $\mathbf{Y} = (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T + \mathbf{E}$ where each column of \mathbf{E} is normal random vector with mean zero and covariance matrix $\boldsymbol{\Sigma}^*$. Assume Assumptions 1-3 hold. Then there exists a constant $C > 0$ such that*

$$\inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*) \in \Theta_K^*} \mathbb{E}_* \left\{ \|\hat{\boldsymbol{\mu}}\hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right\} \geq C (s \log p + n \log K)$$

for sufficiently large n , where \mathbb{E}_* denotes the expected value with respect to $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)$.

The key challenge in the proof of Theorem 1 lies in designing suitable subsets of the parameter space for $(\boldsymbol{\mu}, \mathbf{L})$. We construct three parameter subspaces, each essentially fixing $(\boldsymbol{\mu})_S$, S , and \mathbf{L} , respectively. By controlling the Kullback–Leibler diameter and entropy of each subspace, we apply Fano’s lemma to derive minimax lower bounds on the convergence rate in each subspace, integrating them to obtain the final minimax lower bound.

The minimax lower bound consists of two parts: $s \log p$ and $n \log K$. The $s \log p$ term describes the logarithmic complexity of selecting s non-zero coordinates among p variables. It appears repeatedly in the minimax rates for high-dimensional problems where sparsity plays an important role, including the sparse normal means problem (Castillo and van der Vaart, 2012) and the sparse linear regression (Castillo et al., 2015). The term $n \log K$ comes from the logarithmic complexity of assigning n points into K clusters and also appears in the minimax risk for parameter estimation in stochastic block models (Ghosh et al., 2020).

Remark 2 *This result fills the gap in the literature of high dimensional low-rank matrix estimation, particularly in scenarios where both sparsity and clustering structures exist. When \mathbf{Y} is a $p \times n$ random matrix that can be written as $\mathbf{Y} = \mathbf{X} + \mathbf{E}$, where \mathbf{E} is a $p \times n$ noise matrix whose entries are independent standard normal random variables and \mathbf{X} is a $p \times n$ rank- K matrix, Yang et al. (2016) showed that, if \mathbf{X} not only is low rank but also has only an $s \times l$ non-zero submatrix, then the minimax lower bound is $K(s+l) + s \log(ep/s) + l \log(en/l)$. Our minimax lower bound is sharper than the above bound because the right singular subspace induced by \mathbf{L} contains a clustering structure, whereas the matrix \mathbf{X} considered in Yang et al. (2016) does not have a structured right singular subspace.*

2.3 Minimax upper bound and a constrained maximum likelihood estimator

From the frequentist perspective, an ideal method for parameter estimation in a well-specified statistical model is the maximum likelihood estimator (MLE). In this subsection, we propose a constrained MLE for estimating the mean matrix $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$. We prove that the risk bound of this estimator achieves the minimax lower bound, thereby showing that the minimax lower bound coincides with the minimax risk modulus a multiplicative constant.

Assuming the number of clusters K is known, we consider the parameter space $\Theta_K = \{(\boldsymbol{\mu}, \mathbf{L}) : \boldsymbol{\mu} \in \mathbb{R}^{p \times K}, \mathbf{L} \in \mathcal{L}_K, |\text{supp}(\boldsymbol{\mu})| \leq s\}$ and define the following constrained MLE

$$(\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}) = \arg \min_{(\boldsymbol{\mu}, \mathbf{L}) \in \Theta_K} \|\mathbf{Y} - \boldsymbol{\mu} \mathbf{L}^T\|_F^2. \quad (2)$$

It is worth noting that the parameter space in the constrained MLE is not necessarily compact. However, by characterizing a compact neighbor of $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$ and controlling the complexity inside and outside this neighbor separately, we can establish the risk bound of the constrained MLE and show that it achieves the minimax lower bound in Theorem 1.

Theorem 3 *Suppose that $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*) \in \Theta_K^*$ and $(\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}})$ is defined as in (2) and Assumptions 1-3 hold. Then there exists some constant $c > 0$ such that*

$$\sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*) \in \Theta_K^*} \mathbb{E}_* \left\{ \left\| (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T \right\|_F^2 \right\} \leq c(n \log K + s \log p),$$

where \mathbb{E}_* denotes the expected value with respect to $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)$.

The proof of Theorem 3 relies on transforming the upper bound of the convergence rate into an empirical process on $\tilde{\Theta}_K$, which is the normalization of the feasible set of the estimator Θ_K . This poses challenges due to the infinite entropy of Θ_K . To overcome this, we decompose $\tilde{\Theta}_K$ into layers $(\mathcal{E}_j^K)_{j \in \mathbb{Z}}$ with manageable entropy. We classify these layers into three scenarios: small volume, large volume, and intermediate volume. For the small volume case, we establish that \mathbf{L} equals \mathbf{L}^* up to a permutation within \mathcal{E}_j^K , simplifying the complexity to focus on $\boldsymbol{\mu}$ alone. In the large volume case, we control the entropy separately for the spaces of $\boldsymbol{\mu}$ and \mathbf{L} . Lastly, in the intermediate volume case, we leverage the ellipsoidal nature of \mathcal{E}_j^K and its packing to bound the covering numbers efficiently.

Despite the theoretical optimality of the constrained MLE, it is computationally intractable in general since the feasible set Θ_K is nonconvex and involves discrete structures. In addition, the implementation of the constrained MLE requires to pre-specify the sparsity

level s and the number of clusters K , which are usually unknown in practice. These computational challenges motivate us to develop a Bayesian method that can be implemented conveniently using an MCMC sampler without specifying s and K *a priori*.

2.4 Bayesian sparse high-dimensional Gaussian mixture model

As described in the previous subsection, the optimization-based constrained MLE is computationally intractable due to the non-convexity and discrete structure of the problem. One may apply the EM algorithm, which iterates between a clustering step given the recently updated parameter values and a parameter estimation step given the recently updated cluster memberships, to address this issue. For example, Cai et al. (2019) proposed an approach that estimates the sparse discriminant vector and obtains the clustering memberships in the Expectation step to avoid the singularities of sample covariance matrices in high dimensions. Another approach is spectral clustering (Luxburg, 2004). However, the optimality of spectral clustering is only established when $p = o(n\Delta)$ without sparsity structure (Löffler et al., 2021). In this subsection, we propose a Bayesian approach to estimate the high-dimensional sparse Gaussian mixture model. As will be seen later, the proposed Bayesian method has a minimax-optimal posterior contraction rate.

We deliberately consider a misspecified sampling model: $\mathbf{Y}_i = \boldsymbol{\mu}_{z_i} + \boldsymbol{\epsilon}_i$, where the error term $\boldsymbol{\epsilon}_i$ follows a multivariate normal distribution with mean vector zero and identity covariance matrix \mathbf{I}_p . This intentional simplification is motivated by theoretical convenience, since our primary focus lies in the mean matrix $\boldsymbol{\mu}\mathbf{L}^T$. As we will prove later, as long as the spectrum of the true covariance matrix $\boldsymbol{\Sigma}^*$ is bounded, the posterior distribution of $\boldsymbol{\mu}\mathbf{L}$ concentrates on the true parameter $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$ at a minimax-optimal rate.

To promote sparsity, we use the spike-and-slab LASSO prior (Ročková and George, 2018) for the mean vectors of clusters. The spike-and-slab LASSO prior can be viewed as a continuous relaxation of the spike-and-slab prior (Mitchell and Beauchamp, 1988), which is a mixture of a point mass at zero (referred to as the “spike” distribution) and an absolutely continuous distribution (referred to as the “slab” distribution). Formally, for $\mathbf{x} \in \mathbb{R}^p$, the spike-and-slab LASSO prior is defined as follows: for $j \in [p]$, $\pi(x_j \mid \lambda_0, \lambda_1, \xi_j) = (1 - \xi_j)\psi(x_j \mid \lambda_0) + \xi_j\psi(x_j \mid \lambda_1)$ and $(\xi_j \mid \theta) \sim \text{Bernoulli}(\theta)$, where $\psi(x \mid \lambda) = (\lambda/2) \exp(-\lambda|x|)$ is the density of Laplace distribution with mean 0 and variance $2/\lambda^2$. By assuming $\lambda_0 \gg \lambda_1$, $\psi(x_j \mid \lambda_0)$ closely resembles the “spike” distribution in the spike-and-slab prior since it is highly concentrated at 0, whereas $\psi(x_j \mid \lambda_1)$ plays the role of the “slab” distribution. We follow the notation in Ročková and George (2018) and use $\text{SSL}(\lambda_0, \lambda_1, \theta)$ to denote this prior model. In the context of our proposed sparse Gaussian mixture model, we define the joint-SSL($\lambda_0, \lambda_1, \theta$) as follows to further incorporate the case where the mean vectors $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K \in \mathbb{R}^p$ share the same sparsity pattern: given K , for $j \in [p]$,

$$\begin{aligned} \pi(\boldsymbol{\mu}_{1j}, \dots, \boldsymbol{\mu}_{Kj} \mid \lambda_0, \lambda_1, \xi_j) &= \prod_{k=1}^K ((1 - \xi_j)\psi(\boldsymbol{\mu}_{kj} \mid \lambda_0) + \xi_j\psi(\boldsymbol{\mu}_{kj} \mid \lambda_1)), \\ (\xi_j \mid \theta) &\sim \text{Bernoulli}(\theta). \end{aligned}$$

Under this prior distribution, the random vectors $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K$ are conditionally independent given K and a sparsity indicator vector $\boldsymbol{\xi} \in \{0, 1\}^p$ which controls the common sparsity

structure. We further assume that $\theta \sim \text{Beta}(1, \beta_\theta)$, where $\beta_\theta = p^{1+\kappa} \log p$ for some constant $\kappa > 0$. The choice of the hyperparameter β_θ is selected for technical reasons.

We now specify the sparse Gaussian mixture model. Given K , for cluster membership indicators z_1, \dots, z_n , we assign a categorical prior with a K -dimensional probability vector $\mathbf{w} = (w_1, \dots, w_K)^T$, whose hyperprior distribution is a K -dimensional symmetric Dirichlet distribution with the shape parameter $\alpha > 0$. We assign a joint-SSL prior for the mean vectors $\{\boldsymbol{\mu}_k\}_{k=1}^K$ to adapt to the joint sparsity. To allow for an unknown K , we further assign a truncated Poisson distribution to K by letting $\pi(K) \propto e^{-\lambda} \lambda^K / K!$, $K \in [K_{\max}]$, where K_{\max} is a conservative upper bound for K and should be large enough in practice. Thus, the proposed Bayesian sparse Gaussian mixture model can be expressed as follows:

$$(\mathbf{Y}_1, \dots, \mathbf{Y}_n \mid \mathbf{z}, \boldsymbol{\mu}) \sim N_p(\boldsymbol{\mu}_k, \mathbf{I}_p) \quad \text{independently,} \quad (3)$$

$$(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K \mid K, \theta) \sim \text{joint-SSL}(\lambda_0, \lambda_1, \theta), \quad (4)$$

$$(z_1, \dots, z_n \mid \mathbf{w}, K) \sim \text{Categorical}(\mathbf{w}) \quad \text{independently,} \quad (5)$$

$$(\mathbf{w} \mid K) \sim \text{Dirichlet}_K(\alpha), \quad (6)$$

$$\pi(K) \propto \frac{e^{-\lambda} \lambda^K}{K!}, \quad K \in [K_{\max}] \quad (7)$$

$$\theta \sim \text{Beta}(1, \beta_\theta) \quad \text{where } \beta_\theta = p^{1+\kappa} \log p. \quad (8)$$

The use of sparsity-enforcing priors in Gaussian mixture models has been widely explored in Bayesian literature. For example, Tadesse et al. (2005) and Gao et al. (2020) proposed a discrete subset-selection prior for clustering high-dimensional data. In this paper, we employ a continuous spike-and-slab shrinkage prior. While Tadesse et al. (2005) focused solely on computational algorithms without theoretical analysis, and Gao et al. (2020) primarily investigated the theoretical results of parameter estimation, our main contribution lies in establishing the theoretical properties of the proposed model concerning both parameter estimation and mis-clustering error. This represents the first effort in developing a general theoretical framework for Bayesian analyses of high-dimensional clustering.

3. Theoretical Properties

In this section, we establish the theoretical guarantees of the proposed Bayesian sparse Gaussian mixture model.

3.1 Posterior contraction rate

In this subsection, we show that the posterior contraction rate with respect to the Frobenius norm metric is minimax optimal under the proposed Bayesian sparse Gaussian mixture model. All the proofs are deferred to the Appendix.

By the Bayes formula, the posterior distribution of $\boldsymbol{\mu}$ and \mathbf{L} can be written as

$$\Pi\{(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E} \mid \mathbf{Y}\} = \frac{\int_{\mathcal{E}} p_n(\mathbf{Y} \mid \boldsymbol{\mu}, \mathbf{L}) / p_n(\mathbf{Y} \mid \boldsymbol{\mu}^*, \mathbf{L}^*) \Pi(d\boldsymbol{\mu}d\mathbf{L})}{\int_{\Theta} p_n(\mathbf{Y} \mid \boldsymbol{\mu}, \mathbf{L}) / p_n(\mathbf{Y} \mid \boldsymbol{\mu}^*, \mathbf{L}^*) \Pi(d\boldsymbol{\mu}d\mathbf{L})},$$

where $p_n(\mathbf{Y} \mid \boldsymbol{\mu}, \mathbf{L}) = (2\pi)^{-np/2} \exp(-\|\mathbf{Y} - \boldsymbol{\mu}\mathbf{L}^T\|_F^2/2)$ is the likelihood of the data matrix \mathbf{Y} with identity covariance matrix and \mathcal{E} is any measurable subset of $\Theta = \bigcup_{K=1}^{K_{\max}} \mathbb{R}^{p \times K} \times \mathcal{L}_K$. In Theorem 4, we derive the posterior contraction rate under the proposed Bayesian model.

Theorem 4 *Let \mathbf{Y} be generated from a mixture of K^* Gaussian distributions as in (1) with the true mean vectors $\boldsymbol{\mu}^* = [\boldsymbol{\mu}_1^*, \dots, \boldsymbol{\mu}_{K^*}^*]$ and the true cluster membership matrix \mathbf{L}^* , where $|\text{supp}(\boldsymbol{\mu}^*)| \leq s$. Suppose Assumptions 1 - 3 hold. Let $\boldsymbol{\mu}$ and \mathbf{L} follow the prior specification in (3)-(8) with some hyperparameters $\kappa > 0$, $\alpha \geq 1$, $\lambda_0 \geq 2 \log(p/s) \sqrt{np/(s \log p)}$ and $1/n^\gamma \lesssim \lambda_1 \lesssim \sqrt{s \log p / \|\boldsymbol{\mu}^*\|_F}$ for some constants $\gamma > 0$. Then, we have*

$$\Pi \left\{ (\boldsymbol{\mu}, \mathbf{L}) : \|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \geq M(s \log p + n \log K_{\max}) \middle| \mathbf{Y} \right\} \rightarrow 0$$

in $\mathbb{P}_{(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)}$ -probability, for every large constant M and $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*) \in \Theta_K^*$.

The proof of Theorem 4 adopts a modified “testing-and-prior-concentration” approach (Ghosal et al., 2000), tailored to address the unique challenges arising from model misspecification in our setting. We rigorously prove the three conditions: (1) The prior distribution puts a sufficient mass on the neighbourhood of the true parameter $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$; (2) There exists a test function which can distinguish $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$ from the complement of its neighbourhood in a subset of the parameter space; (3) The prior puts almost all mass on the subset of parameter space in condition (2).

Remark 5 *Recall from Section 2 that the minimax lower bound contains the true number of clusters K^* , which is unknown in many applications. The posterior contraction rate obtained in Theorem 4 contains a logarithmic factor of the upper bound K_{\max} for K^* . If we further assume that $K_{\max} \asymp (K^*)^q$ for some constant $q \geq 1$, the posterior contraction rate matches the minimax lower bound in Theorem 1 and is optimal thereafter. For λ_1 in the joint-SSL prior, if we further assume $\|\boldsymbol{\mu}_k^*\|_2^2 = O(s)$ for any $k \in [K^*]$, then the upper bound of λ_1 can be relaxed to $\lambda_1 \lesssim \sqrt{\log p / K_{\max}}$, which is a mild condition and can be easily satisfied in practice.*

We assume that the cluster mean vectors are jointly sparse. However, Theorem 4 can be easily generalized to the case where the cluster centers do not share the common sparsity. Specifically, each mean vector $\boldsymbol{\mu}_k$ has at most s non-zero coordinates but the indices of the non-zero coordinates are not necessarily the same across $k \in [K]$. Clearly, the matrix $\boldsymbol{\mu} = [\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K]$ is jointly Ks -sparse. To adapt to the column-wise sparsity of $\boldsymbol{\mu}$, we modify the prior model by letting $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K$ follow the SSL prior independently given K :

$$(\boldsymbol{\mu}_k | K, \theta) \sim \text{SSL}(\lambda_0, \lambda_1, \theta) \quad \text{for } k = 1, \dots, K. \quad (6')$$

The following corollary gives the posterior contraction rate under such a modification.

Corollary 6 *Let \mathbf{Y} be generated from a mixture of K^* Gaussian distributions as in (1) with the true mean vectors $\boldsymbol{\mu}^* = [\boldsymbol{\mu}_1^*, \dots, \boldsymbol{\mu}_{K^*}^*]^T$ and the true cluster membership matrix \mathbf{L}^* , where $|\text{supp}(\boldsymbol{\mu}^*)| \leq K^*s$. Suppose Assumptions 1-3 hold. Let $\boldsymbol{\mu}$ and \mathbf{L} follow the prior specification in (3), (6'), (5)-(8) with the same hyperparameters as in Theorem 4. Then, we have*

$$\Pi \left\{ (\boldsymbol{\mu}, \mathbf{L}) : \|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \geq M(sK_{\max} \log p + n \log K_{\max}) \middle| \mathbf{Y} \right\} \rightarrow 0$$

in $\mathbb{P}_{(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)}$ -probability, for every large constant M and $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*) \in \Theta_K^*$.

Remark 7 *Corollary 6 can be easily extended to the case when the mean vectors have not only different sparsity structures, but also distinct sparsity sizes, i.e., $|\text{supp}(\boldsymbol{\mu}_k)| \neq |\text{supp}(\boldsymbol{\mu}_{k'})|$ for some $k \neq k'$. In such a case, the first term of the posterior contraction rate becomes to $\sum_{k=1}^{K^*} |\text{supp}(\boldsymbol{\mu}_k)| K_{\max} \log p$.*

3.2 Mis-clustering error

Recovering cluster memberships is always a focal objective for clustering problems. In this subsection, we obtain a mis-clustering error bound of the proposed Bayesian model *a posteriori* based on the posterior contraction result for parameter estimation in Theorem 4. For any two cluster membership vectors $\mathbf{z}, \mathbf{z}' \in ([K])^n$, define the minimum Hamming distance $d_H(\mathbf{z}, \mathbf{z}') = (1/n) \inf_{\tau \in S_K} \sum_{i=1}^n \mathbb{1}\{z_i \neq \tau(z'_i)\}$ as the mis-clustering rate between \mathbf{z} and \mathbf{z}' , where S_K is the set of all permutations on $[K]$. Let $\sigma_{\max}(\mathbf{X})$ and $\sigma_{\min}(\mathbf{X})$ denote the largest and smallest non-zero singular value of matrix \mathbf{X} , respectively. Below, we obtain the posterior contraction result for the mis-clustering error measured by d_H .

Theorem 8 *Assume the conditions in Theorem 4 hold and $n_k^* = \sum_{i=1}^n \mathbb{1}(z_i^* = k) \rightarrow \infty$ for all k . Let $n_{\min}^* \triangleq \min_{k \in [K]} n_k^*$ and $n_{\max}^* \triangleq \max_{k \in [K]} n_k^*$. Then we have*

$$\Pi \left\{ nd_H(\mathbf{z}, \mathbf{z}^*) \geq \frac{M(n_{\max}^*)^3 \sigma_{\max}(\boldsymbol{\mu}^*)^2}{(n_{\min}^*)^4 \sigma_{\min}(\boldsymbol{\mu}^*)^4} (s \log p + n \log K_{\max}) \middle| \mathbf{Y} \right\} \rightarrow 0$$

in $\mathbb{P}_{(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)}$ -probability for every large constant M .

The main challenge in proving Theorem 8 is translating clustering accuracy into errors regarding the right singular subspace of $\boldsymbol{\mu} \mathbf{L}^T$. By leveraging a variant of the Davis-Kahan Theorem, we bound the distance of the right singular subspace of $\boldsymbol{\mu} \mathbf{L}^T$ using the distance of $\boldsymbol{\mu} \mathbf{L}^T$. Through geometric analysis, we show that correct clustering is achieved if the errors of right singular vectors are sufficiently small. This allows us to control mis-clustering error by the error of $\boldsymbol{\mu} \mathbf{L}^T$, and the desired result follows by directly applying Theorem 4.

Remark 9 *If we assume $(s \log p)/(n \|\boldsymbol{\mu}^*\|_F^2) \rightarrow 0$, then by Theorem 8, the proportion of the mis-clustered data points is asymptotically negligible with a high posterior probability provided that $(n_{\max}^*)^3 \sigma_{\max}(\boldsymbol{\mu}^*)^4 = O((n_{\min}^*)^4 \sigma_{\min}(\boldsymbol{\mu}^*)^4)$ as $n \rightarrow \infty$. Moreover, if we further assume that $\sigma_{\max}(\boldsymbol{\mu}^*) \lesssim \sigma_{\min}(\boldsymbol{\mu}^*)$ and $n_{\max}^* \lesssim n_{\min}^* = O(n)$ (which means that the sizes of the smallest cluster and the largest cluster are of the same order as n), then the number of mis-clustered data points, i.e., $nd_H(\mathbf{z}, \mathbf{z}^*)$, is asymptotically bounded by a constant with a high posterior probability because $\|\boldsymbol{\mu}^*\|_F^2 \asymp \sigma_{\min}(\boldsymbol{\mu}^*)^2$ in this case.*

Remark 10 *Azizyan et al. (2013) and Cai et al. (2019) also studied high-dimensional clustering with the sparsity assumption. However, they only considered the case when the number of clusters was 2. Assuming that $\boldsymbol{\mu}_1^* - \boldsymbol{\mu}_2^*$ was sparse, Azizyan et al. (2013) showed that the minimax optimal convergence rate of mis-clustering was $\sqrt{s \log p/n}/\Delta^2$ when the two clusters had same mixing weights and isotropic covariance matrices. Assuming that the discriminant direction vector $\beta^* = (\boldsymbol{\Sigma}^*)^{-1}(\boldsymbol{\mu}_1^* - \boldsymbol{\mu}_2^*)$ was sparse, Cai et al. (2019) showed that the convergence rate of the excess mis-clustering error, defined as the difference between the mis-clustering error and the optimal mis-classification error obtained by Fisher's*

linear discriminant rule when cluster-specific parameters were known, achieved the minimax optimal rate of $s \log p/n$. However, the convergence rate of mis-clustering error was not investigated. In addition, Li et al. (2017), Azizyan et al. (2013) and Cai et al. (2019) focused on the two-cluster problem, but the minimax optimal result for high-dimensional sparse clustering with $K > 2$ clusters was not studied. In contrast, we allow K to grow moderately with the sample size n .

4. Simulation Studies

We evaluate the empirical performance of the proposed Bayesian method for sparse Gaussian mixtures through analyses of synthetic data sets. Posterior inference is carried out through an MCMC sampler, the details of which are provided in Appendix C. We also compare the performance of our model with four competitors: principal component analysis K-means (PCA-KM), sparse K-means (SKM) (Witten and Tibshirani, 2010), Clustering of High-dimensional Gaussian Mixtures with the EM (CHIME), and Gaussian-mixture-model-based clustering (MClust) (Fraley et al., 2012). In particular, PCA-KM is a two-stage approach that first performs a PCA to reduce dimensionality and then applies a K-means algorithm to the principal components. SKM is a generalization of the K-means in high dimensions to find clusters and important features (i.e., the non-zero coordinates) simultaneously. CHIME is a high-dimensional clustering approach based on an EM algorithm. To overcome the issue that the sample covariance matrix may not be invertible and thus the subsequent estimate of \mathbf{z} is not available, CHIME focuses on the so-called sparse discriminant vector and directly use it in the Fisher discriminant rule to estimate cluster memberships. Note that the performance of CHIME is quite sensitive to the choice of initial values. Throughout simulation examples in this section, we set the initial values of CHIME to be the output of K-means. For PCA-KM and SKM, we choose the number of clusters via Silhouette method (Rousseeuw, 1987), with the range of K being from 2 to 10. For MClust and CHIME, the number of clusters is estimated via Bayesian information criterion (BIC).

4.1 Simulation setup

We consider three simulation scenarios. Scenario I is designed to evaluate the proposed Bayesian method in terms of clustering accuracy with varying numbers of clusters and support sizes of the mean vectors. The data matrix \mathbf{Y} is of size $p \times n$ with $p = 400$ and $n = 200$. We assume that the true number of clusters K^* ranges over $\{3, 5\}$ and the support size s ranges over $\{6, 12\}$. We use \mathcal{S} to denote the set of non-zero coordinates and let the first s coordinates of the cluster means be non-zero, i.e., $\mathcal{S} = \{1, 2, \dots, s\}$. For each $K^* \in \{3, 5\}$, the true cluster assignment z_i^* is generated from a categorical distribution: $z_i^* \sim \text{Cat}(\mathbf{p}_{K^*})$, where $\mathbf{p}_3 = (0.3, 0.3, 0.4)$ and $\mathbf{p}_5 = (0.2, 0.2, 0.2, 0.2, 0.2)$. When $K^* = 3$, the three cluster mean vectors are $(\boldsymbol{\mu}_1^*)_{\mathcal{S}} = 3 \times (1, 1, \dots, 1)^T$, $(\boldsymbol{\mu}_2^*)_{\mathcal{S}} = -1.5 \times (1, 1, \dots, 1)^T$ and $(\boldsymbol{\mu}_3^*)_{\mathcal{S}} = (0, \dots, 0)^T$, where $(\boldsymbol{\mu}_1^*)_{\mathcal{S}}, (\boldsymbol{\mu}_2^*)_{\mathcal{S}}, (\boldsymbol{\mu}_3^*)_{\mathcal{S}} \in \mathbb{R}^s$. For $K^* = 5$, the five cluster mean vectors are $(\boldsymbol{\mu}_1^*)_{\mathcal{S}} = 4 \times (1, 1, \dots, 1)^T$, $(\boldsymbol{\mu}_2^*)_{\mathcal{S}} = -4 \times (1, 1, \dots, 1)^T$, $(\boldsymbol{\mu}_3^*)_{\mathcal{S}} = (0, \dots, 0)^T$, $(\boldsymbol{\mu}_4^*)_{\mathcal{S}} = 4 \times (-1, 1, -1, 1, \dots, -1, 1)^T$ and $(\boldsymbol{\mu}_5^*)_{\mathcal{S}} = 1.5 \times (1, -1, 1, -1, \dots, 1, -1)^T$. Given K^* and z_i^* 's, data are generated from $\mathbf{Y}_i \sim N(\boldsymbol{\mu}_{z_i^*}^*, \mathbf{I}_p)$.

Scenario II focuses on the case when small clusters exist. The data matrix \mathbf{Y} consists of $n = 200$ observations with dimension $p = 400$. We assume that the true number of clusters

$K^* = 3$ and the support size $s = 8$. Similarly as Scenario I, we set $\mathcal{S} = \{1, 2, \dots, s\}$. The mean vectors over the support \mathcal{S} in the three clusters are $(\boldsymbol{\mu}_1^*)_{\mathcal{S}} = (5, 2, \dots, 5, 2)^T$, $(\boldsymbol{\mu}_2^*)_{\mathcal{S}} = (10, 5, \dots, 10, 5)^T$, and $(\boldsymbol{\mu}_3^*)_{\mathcal{S}} = (15, 2, \dots, 15, 2)^T$, respectively. For each observation i , its simulated true cluster assignment z_i^* is generated from a categorical distribution independently: $z_i^* \sim \text{Cat}(0.02, 0.48, 0.5)$. Given K^* and z_i^* 's, data are generated from $\mathbf{Y}_i \sim N(\boldsymbol{\mu}_{z_i^*}^*, \boldsymbol{\Sigma}_{z_i^*})$, where $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_3 = \mathbf{I}_p$, and $\boldsymbol{\Sigma}_2$ is a diagonal matrix whose diagonal entries equal 4 on the coordinates in the support \mathcal{S} and 1 elsewhere.

Scenario III aims to investigate the robustness of the proposed Bayesian method to the misspecification of the sampling distribution. The true distribution of the data is assumed to be a mixture of multivariate t distributions but we use the Gaussian mixtures as the working likelihood. The data matrix \mathbf{Y} consists of $n = 200$ observations of multivariate t -mixtures with dimension $p = 400$ and a degree of freedom 5. The number of clusters is set to $K^* = 3$, and the first $s = 8$ coordinates of cluster mean vectors are non-zero. We generate cluster assignments z_i^* 's from a categorical distribution: $z_i^* \sim \text{Cat}(0.2, 0.4, 0.4)$ independently for $i \in [n]$, and let the cluster mean vectors $\boldsymbol{\mu}_1^*, \boldsymbol{\mu}_2^*, \boldsymbol{\mu}_3^*$ and the covariance matrices $\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2, \boldsymbol{\Sigma}_3$ be the same as those in Scenario II for each multivariate t -cluster.

For each of the three scenarios, we apply the proposed Bayesian sparse Gaussian mixture model to the simulated data with 100 repeated simulations. In each simulation, we compute posterior inference using the developed MCMC sampler with 1000 burn-in iterations and another 4000 iterations for post-burn-in samples. The upper bound of the number of clusters is set to be $K_{\max} = 20$. We set the hyperparameters κ , λ_0 , and λ_1 in the spike-and-slab LASSO prior to be 0.1, 100, and 1 respectively, and λ in the truncated Poisson prior for K to be 2. The estimated number of clusters and cluster assignments under the proposed Bayesian method are reported based on the posterior mode of z_i 's from post-burn-in MCMC samples. The proposed Bayesian method, Mclust, PCA-KM, and SKM are performed under R with version 4.2.1 and CHIME is performed under Matlab with version 9.11 (R2021b).

4.2 Simulation results

We first investigate the performance of the proposed Bayesian sparse Gaussian mixture model against the four competitors in Scenario I. We focus on the following three objectives: identification of the number of clusters, the clustering accuracy, and the cluster-wise mean vector estimation accuracy. The proposed Bayesian method can successfully recover the simulated true number of clusters. Specifically, when $K^* = 3$, the proposed method identifies 3 clusters in 85 replicates out of 100 replicates for $s = 6$ and in 98 replicates for $s = 12$; when $K^* = 5$, the proposed method identifies 5 clusters in 83 replicates out of 100 replicates for $s = 6$ and in 98 replicates for $s = 12$. In contrast, all the four competitors underestimate the number of clusters. In particular, when $K^* = 3$, the estimated number of clusters using the four competitors all equal to 2 in 100 simulation replicates. When $K^* = 5$, PCA-KM, SKM, Mclust, and CHIME only correctly estimate the number of clusters in 6, 0, 4, and 3 out of 100 replicates for $s = 6$, and in 8, 0, 9, and 3 out of 100 replicates for $s = 12$. Figure 1 and Appendix Figure A1 plot the simulated true cluster memberships and the estimated clustering results under the proposed Bayesian method and the four competitors for one randomly selected simulation replicate when $K^* = 3$, $s = 6$, and $K^* = 5$, $s = 6$, respectively. We can see that the four competitors cannot well distinguish clusters with a

certain degree of overlapping, e.g., the green and blue clusters in the upper left panel of Figure 1, while the proposed Bayesian method can successfully separate them.

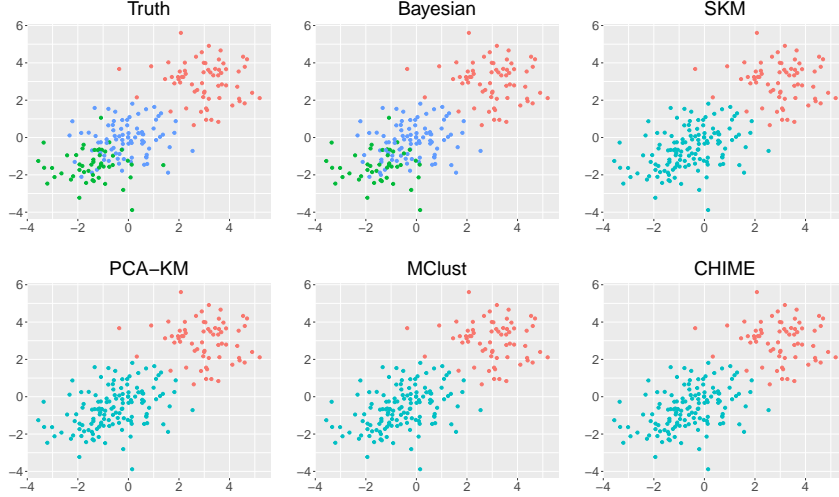


Figure 1: Clustering results of different methods compared to the true cluster memberships in Scenario I with $K^* = 3$ and $s = 6$ in a randomly selected simulation replicate. Data points are projected onto the subspace of the first two coordinates and different colors correspond to different estimated cluster memberships.

In terms of clustering accuracy, we use the adjusted Rand index (ARI) (Rand, 1971) as the evaluation metric. Specifically, let \mathcal{C}^* and \mathcal{C}^\dagger denote the true and estimated partitions of $[n]$ respectively, and let $K^* \triangleq |\mathcal{C}^*|$, $K^\dagger \triangleq |\mathcal{C}^\dagger|$. Denote n_k^* the size of k th cluster in \mathcal{C}^* and $n_{k'}^\dagger$ as the size of k' th cluster in \mathcal{C}^\dagger . Let $n_{k,k'}$ be the number of observations that are assigned to both the k th cluster in \mathcal{C}^* and k' th cluster in \mathcal{C}^\dagger . Then the ARI is defined as

$$\text{ARI}(\mathcal{C}^*, \mathcal{C}^\dagger) = \frac{\sum_{k=1}^{K^*} \sum_{k'=1}^{K^\dagger} \binom{n_{k,k'}}{2} - \sum_{k=1}^{K^*} \binom{n_k^*}{2} \sum_{k'=1}^{K^\dagger} \binom{n_{k'}}{2} / \binom{n}{2}}{\left(\sum_{k=1}^{K^*} \binom{n_k^*}{2} + \sum_{k'=1}^{K^\dagger} \binom{n_{k'}}{2} \right) / 2 - \sum_{k=1}^{K^*} \binom{n_k^*}{2} \sum_{k'=1}^{K^\dagger} \binom{n_{k'}}{2} / \binom{n}{2}}.$$

Table 1 reports the average ARIs of the clustering results of the proposed Bayesian method against the four competitors across 100 simulation replicates under Scenario I. The proposed Bayesian method outperforms the four alternatives in terms of higher ARIs in all settings.

We then examine the cluster-wise mean vector estimation accuracy by computing $\|\hat{\boldsymbol{\mu}}_1 - \boldsymbol{\mu}_1^*\|_2$ under the proposed Bayesian method and alternatives, where $\hat{\boldsymbol{\mu}}_1$ is the estimated mean vector under different methods. Specifically, $\hat{\boldsymbol{\mu}}_1$ under the proposed Bayesian method is the posterior mean of $\boldsymbol{\mu}_1$. CHIME and MClust directly return $\hat{\boldsymbol{\mu}}_1$ since they are model-based methods. For PCA-KM and SKM, we use the empirical means induced from their estimated clustering memberships as $\hat{\boldsymbol{\mu}}_1$ since they are based on K-means. Figure 2 presents the boxplots of $\|\hat{\boldsymbol{\mu}}_1 - \boldsymbol{\mu}_1^*\|_2$ when $K^* = 3$ and $s = 12$ across 100 simulation replicates under different methods, showing that the proposed Bayesian method yields the smallest error

	$K^* = 3$		$K^* = 5$	
	$s = 6$	$s = 12$	$s = 6$	$s = 12$
Bayesian	0.84 (0.19)	0.98 (0.01)	0.94 (0.03)	0.99 (0.01)
PCA-KM	0.54 (0.04)	0.55 (0.04)	0.64 (0.17)	0.60 (0.18)
MClust	0.54 (0.04)	0.55 (0.04)	0.81 (0.13)	0.78 (0.05)
SKM	0.55 (0.04)	0.55 (0.04)	0.54 (0.21)	0.74 (0.13)
CHIME	0.53 (0.10)	0.63 (0.18)	0.52 (0.27)	0.54 (0.29)

Table 1: Average (standard deviation) adjusted Rand indices (ARIs) in Scenario I with different choice of the number of clusters K^* and support size s .

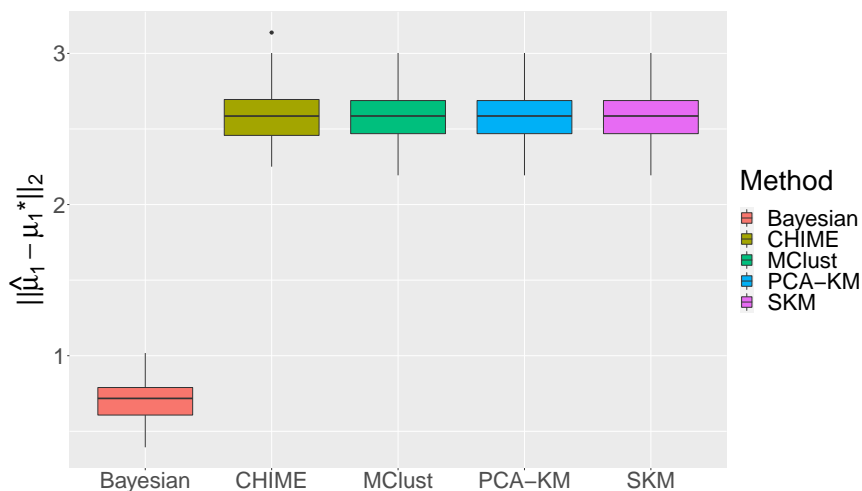


Figure 2: Numerical results of $\|\hat{\mu}_1 - \mu_1^*\|_2$ of different methods in Scenario I with $K^* = 3$ and $s = 12$ across 100 simulation replicates.

of estimating μ_1^* . Furthermore, Appendix Figure A2 plots the estimated μ_1^* under different methods in one randomly selected simulation replicate, indicating that the proposed Bayesian method recovers the sparsity pattern better than the four competitors. Lastly, we report the running times of all methods in Scenario I in Appendix Table A1.

Scenario II is designed to evaluate the proposed Bayesian method when small clusters exist. Figure 3(a) shows the true clustering assignments in one randomly selected replicate, in which the small cluster (in red color) only contains four data points. Our Bayesian method successfully discovers the small cluster and yields the estimated number of clusters $K = 3$ in 96 out of 100 simulation replicates, resulting in an average ARI of 0.99. Figure 3(b) shows the estimated clustering memberships under the proposed Bayesian method in the same simulation replicate, exactly matching the truth shown in Figure 3(a). In contrast, all four competitors are not able to identify the small cluster and report the estimated number of clusters $K = 2$ in all 100 simulation replicates. We further examine the performance of the

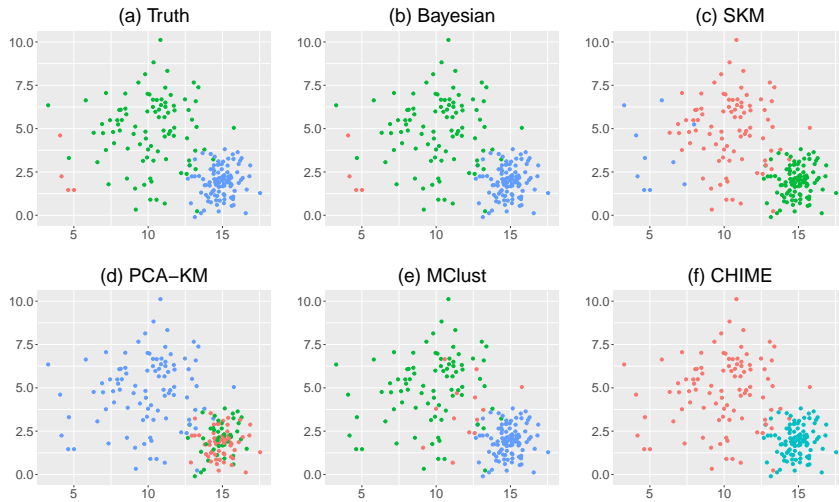


Figure 3: Clustering results of different methods compared with the truth in one randomly selected simulation replicate of Scenario II. (a) Simulated truth. (b) Clustering result under the proposed Bayesian method. (c-f) Clustering results under SKM, PCA-KM, MClust, and CHIME when the number of clusters is fixed to be truth. Data points are projected onto the subspace of the first two coordinates and different colors correspond to different estimated cluster memberships.

four competitors when we pre-specify the number of clusters to be the truth $K = 3$. Figure 3(c) - (f) plot the estimated clustering memberships obtained from the four competitors. We can see that SKM, PCA-KM, and MClust incline to return clusters with relatively balance sizes, leading to inaccurate clustering assignments with the average ARIs being 0.76, 0.84, and 0.79, respectively, across 100 simulation replicates. CHIME only returns two clusters even though we set the number of clusters to be 3, as shown in Figure 3(f).

For Scenario III, the histograms of the estimated number of clusters under different methods are presented in Appendix Figure A3. Figure 4 visualizes the clustering results under different methods in a randomly selected simulation replicate. The proposed Bayesian method successfully identifies three clusters in 76 out of 100 simulation replicates, with an average ARI of 0.97 across 100 simulation replicates. Note that the multivariate t -distribution is heavy-tailed. Since we mis-specify the working model as Gaussian mixtures, it is reasonable to treat some observations as “outliers”, as shown in the upper left panel of Figure 4. Therefore, the proposed Bayesian method tends to assign these “outliers” to small clusters when it overestimates the number of clusters. PCA-KM returns 10 clusters in 59 out of 100 simulation replicates with the average ARI being 0.84. The estimated numbers of clusters of SKM are all 2 in 100 simulation replicates, and the average ARI is 0.52. For model-based methods, i.e., MClust and CHIME, which also use Gaussian mixtures as the working likelihood, their performance are much worse than others as they only identify one cluster in 93 out of 100 replicates, resulting in the average ARIs less than 0.05.




Figure 4: Clustering results under different methods in one randomly selected replicate of Scenario III. Data points are projected onto the subspace of the first two coordinates and different colors correspond to different estimated cluster memberships.

5. Single-cell Sequencing Data Analysis

Recent advances in high-throughput single-cell RNA sequencing (scRNA-Seq) technologies greatly enhance our understanding of cell-to-cell heterogeneity and cell lineages trajectories in development (Cao et al., 2019). One important goal of analyzing scRNA-Seq data is to cluster cells to identify cell subpopulations with different functions and gene expression patterns. The large number of detected genes in scRNA-Seq data makes the task of clustering cells a high-dimensional problem. In this section, we evaluate the proposed Bayesian sparse Gaussian mixture model using a benchmark scRNA-Seq data set (Darmas et al., 2015), which is available at the data repository Gene Expression Omnibus (GSE67835, <https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE67835>). After excluding hybrid cells and filtering out lowly expressed genes (i.e., the total number of RNA-Seq counts over all non-hybrid cells is less than or equal to 10), we have $p = 18568$ genes and $n = 420$ cells in 8 cell types including fetal quiescent cells (110 cells), fetal replicating cells (25 cells), astrocytes cells (62 cells), neuron cells (131 cells), endothelial (20 cells), oligodendrocyte cells (38 cells), microglia cells (16 cells), and OPCs (16 cells). The original count data y_{ji} for gene j in cell i is transformed into continuous type by taking base-2 logarithm and pseudo count 1, i.e., $\log_2(y_{ji} + 1)$. Then we divide each y_{ji} by the total expression of each cell, i.e., $\log_2(y_{ji} + 1) / \sum_{j=1}^p \log_2(y_{ji} + 1)$. Lastly we normalize the data such that the standardized expression levels have zero mean and unit variance for each gene. Note that after normalization the data are no longer sparse and exhibit a skewed distribution; however, most values remain small in magnitude (see Figure 5 for two randomly selected dimensions). In our theoretical studies, the sparsity assumption applies only to the underlying mean vector—not to the raw or normalized data—since the framework explicitly accounts for noise in each dimension. This assumption is primarily introduced to guaran-

tee the method’s theoretical properties, although the algorithm itself remains applicable regardless of the exact sparsity level.




Figure 5: Histograms showing the distribution of normalized expression values for two representative dimensions. The highly skewed distributions demonstrate that most values are concentrated near zero, even after normalization.

We apply the proposed Bayesian method to the scRNA-Seq data with the same hyperparameters as in the simulation study. We run the MCMC sampler for 10000 iterations and discard the first 5000 iterations as burn-in. For comparison, we implement several alternatives, including K-means (KM), PCA-KM, MClust, SKM, K-means after non-negative matrix factorization (NMF-KM) (Zhu et al., 2017), and K-means after t-distributed stochastic neighbor embedding algorithm (tSNE-KM) (Linderman et al., 2019). For PCA-KM and NMF-KM, we first project the data onto the top 10-dimensional feature space, then apply the KM algorithm to cluster the cells. For the KM-based method, the optimal number of clusters is determined by the Silhouette method.

To assess the sparsity structure in this dataset, we first examine ξ , the indicator variable in our spike-and-slab prior that determines whether each gene is associated with the spike or slab component. Our analysis reveals that the posterior means of ξ are predominantly zero, with 18,321 out of 18,568 genes exhibiting a posterior mean of exactly zero. This indicates that a vast majority of genes are strongly associated with the spike component, reflecting a high level of sparsity in the underlying mean structure. Therefore, our model effectively captures the sparsity in the data, even after normalization.

To measure the performance of clustering results, we use not only the aforementioned ARI but another commonly-used criteria in the single-cell literature: normalized mutual information (NMI) (Ghosh and Acharya, 2011). Formally, with the same notations as in Section 4.2, the NMI is defined as

$$\text{NMI}(\mathcal{C}^*, \mathcal{C}^\dagger) = \frac{\sum_{k=1}^{K^*} \sum_{k'=1}^{K^\dagger} \frac{n_{k,k'}}{n} \log \left(\frac{n_{k,k'}}{n} \right) - \sum_{k=1}^{K^*} \frac{n_k^*}{n} - \sum_{k'=1}^{K^\dagger} \frac{n_{k'}^\dagger}{n} \log \left(\frac{n_{k'}^\dagger}{n} \right)}{\sqrt{\sum_{k=1}^{K^*} \frac{n_k^*}{n} \log \left(\frac{n_k^*}{n} \right) \sum_{k'=1}^{K^\dagger} \frac{n_{k'}^\dagger}{n} \log \left(\frac{n_{k'}^\dagger}{n} \right)}}.$$

Methods	Bayesian	SKM	PCA-KM	MClust	NMF-KM	tSNE-KM	KM
Estimate of K	7	3	4	4	6	4	9
ARI	0.84	0.15	0.59	0.33	0.63	0.78	0.79
NMI	0.80	0.22	0.58	0.35	0.61	0.70	0.77

Table 2: Estimate of the number of clusters K , ARIs, and NMIs obtained by applying different methods to the scRNA-Seq data set.

Table 2 reports the ARIs and NMIs under the proposed method and alternatives, showing that the proposed model results in the highest ARI and NMI and achieves the best clustering accuracy. Figure 6 plots the true cell types and the estimated cluster memberships under all the methods. Although the proposed Bayesian method underestimates the number of cell types by 1 and yields $K = 7$, it can identify most cell types except for fetal quiescent and fetal replicating cells. The KM method correctly estimates $K = 8$ cell types. However, it cannot recognize OPC cells and gives two additional small clusters that are not interpretable. Other methods tend to underestimate K . In particular, SKM and MClust estimate $K = 3$ and $K = 4$ respectively, and perform worse than others in terms of much lower ARIs and NMIs, as shown in Table 2. Both PCA-KM and tSNE-KM estimate $K = 4$ by correctly identifying the astrocytes cell type and merging fetal quiescent and fetal replicating cell types into one cluster. For the other five cell types, PCA-KM identifies microglia cell type and merges oligodendrocytes, OPC, endothelial, and neuron cell types into one cluster, while tSNE-KM identifies the neuron cell type and merges oligodendrocytes, OPC, microglia, and endothelial cells as one cluster. NMF-KM is able to identify neuron, fetal quiescent, and fetal replicating cell types but cannot distinguish others. tSNE-KM identifies oligodendrocytes, OPC, microglia, and endothelial cells as one cluster.

We further examine the alternative methods when the number of clusters is set to be the true number of cell types ($K = 8$). Appendix Figure A4 plots the estimated clustering memberships under all the alternative methods. As shown in Figure A4, MClust cannot distinguish fetal quiescent and fetal replicating cell types and merges OPC and oligodendrocytes cell types into one cluster. PCA-KM and tSNE-KM return clusters with relatively similar sizes and hence their performance on small clusters are not satisfactory. SKM and NMF-KM perform significantly worse than others since they do not correctly identify any single cell type. Appendix Table A2 reports the resulting ARIs and NMIs, showing that the proposed Bayesian method still yields the highest ARI and NMI even though the number of clusters is correctly pre-specified for all alternative methods.

6. Discussion

We propose a Bayesian approach for high-dimensional Gaussian mixtures where the cluster mean vectors exhibit certain sparsity structure. We fully investigate the minimax risk for estimating the mean matrix, show that the posterior contraction rate is minimax optimal, and obtain an error bound for the mis-clustering error. Our approach demonstrates superior performance in both simulations and real-world applications.

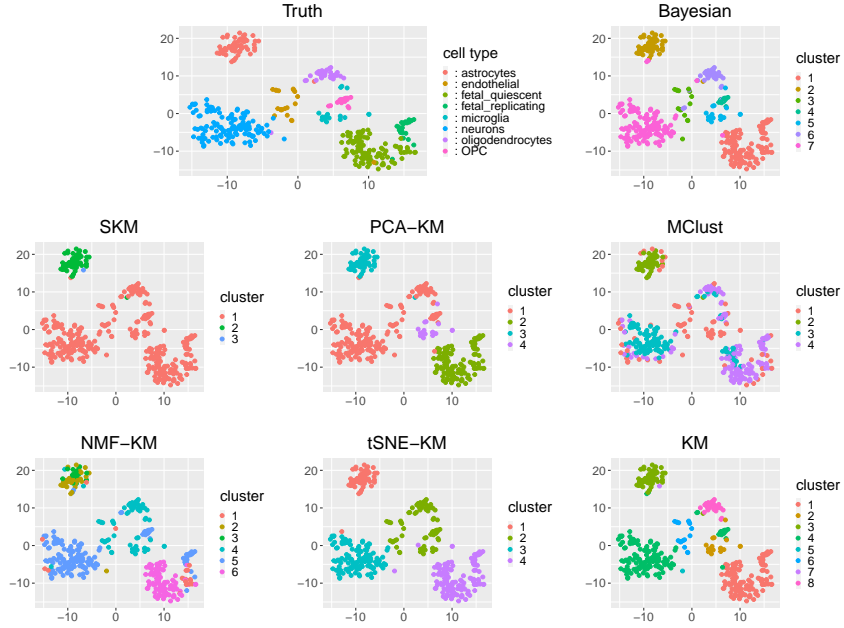


Figure 6: Clustering results of scRNA-Seq data corresponding to different methods. Data points are embedded into two-dimensional subspace by tSNE embedding.

There still exist challenges that need further research. One extension is to consider scenarios where the cluster-specific covariance matrices have some structures, such as sparse spiked structures (Xie et al., 2022). Exploring the theoretical properties of covariance matrix estimation could be an interesting future direction. On the implementation side, algorithms based on Markov chain Monte Carlo can be computationally expensive in ultra-high dimensions. Certain optimization-based alternatives, such as variational Bayes methods (Ray and Szabó, 2021) can be attractive. Developing the underlying backbone theory for variational Bayes approaches can be a promising future research direction as well.

Acknowledgement

The work of Xu was supported by NSF 1918854, NSF 1940107, and NIH R01MH128085.

Appendix A. Proofs for Section 2

Proof of Theorem 1 We define two subspaces of Θ_K^* as follows. We consider the case where $\Sigma^* = I_p$. Without loss of generality, we assume that n/K is an integer. If not, let $n' = \lfloor n/K \rfloor K$. Then, we derive a lower bound for a smaller parameter space $\Theta' \subset \mathbb{R}^{p \times K} \times \mathbb{R}^{n' \times K}$, which also serves as a lower bound for the original parameter space $\Theta_K^* \subset \mathbb{R}^{p \times K} \times \mathbb{R}^{n \times K}$.

We define the first subspace

$$\Theta_{K1}^* = \left\{ (\boldsymbol{\mu}, \mathbf{L}, \boldsymbol{\Sigma}) \in \Theta_K^* : \boldsymbol{\mu} = [\mu_0 \mathbf{w}_1, \dots, \mu_0 \mathbf{w}_K], \mu_0 = \sqrt{\frac{c_1 \log K}{s}}, \text{supp}(\boldsymbol{\mu}) = S, \boldsymbol{\Sigma}^* = \mathbf{I}_p \right\}$$

for some fixed sparsity support $S \subset [p]$ with $|S| = s$, a constant $c_1 > 0$, and $\mathbf{w}_1, \dots, \mathbf{w}_K \in \{0, 1\}^p$. Specifically, we choose the vectors $\mathbf{w}_1, \dots, \mathbf{w}_K$ such that $\|(\mathbf{w}_i)_S\|_0 = s$ for all $i \in [K]$ and $\|(\mathbf{w}_i)_S - (\mathbf{w}_j)_S\|_2^2 > s/2$ for $i \neq j$. By Lemma 4.10 in Massart (2007), there exists $\{(\mathbf{w}_1)_S, \dots, (\mathbf{w}_K)_S\} \subset \{0, 1\}^s$ satisfying these properties for $K \leq s$. Next we define the second subspace. Note that each \mathbf{L} can be associated with a mapping $z : [n] \rightarrow [K]$ such that $\mathbf{l}_i = \mathbf{e}_{z_i}$ where \mathbf{e}_i is the vector whose i th entry is 1 and 0 elsewhere. Then define

$$\Theta_{K2}^* = \left\{ (\boldsymbol{\mu}, \mathbf{L}, \boldsymbol{\Sigma}) \in \Theta_K^* : z^{-1}(k) = \left\{ \frac{(k-1)n}{K} + 1, \dots, \frac{kn}{K} \right\}, \boldsymbol{\Sigma}^* = \mathbf{I}_p \right\}, |\text{supp}(\boldsymbol{\mu})| = s$$

where $\text{supp}(\cdot)$ denotes the set of indices of the non-zero rows of a matrix.

■ We first consider the minimax lower bound over Θ_{K1}^* . By Lemma 4.10 in Massart (2007) we know there exist $\{\mathbf{w}_1, \dots, \mathbf{w}_K\} \subset \{0, 1\}^s$ such that $\|\mathbf{w}_i - \mathbf{w}_j\|^2 > s/2$. Consider an ϵ_1 -ball of Θ_{K1}^* with respect to the metric $d_1((\boldsymbol{\mu}, \mathbf{L}_1), (\boldsymbol{\mu}, \mathbf{L}_2)) = \|\boldsymbol{\mu} \mathbf{L}_1^T - \boldsymbol{\mu} \mathbf{L}_2^T\|_F / \sqrt{n}$. Suppose \mathbf{L}_1 and \mathbf{L}_2 are associated with mappings $z_1, z_2 : [n] \mapsto [K]$. We have

$$4n\epsilon_1^2 > \|\boldsymbol{\mu} \mathbf{L}_1^T - \boldsymbol{\mu} \mathbf{L}_2^T\|_F^2 \geq \mu_0^2 \sum_{i=1}^n \|\mathbf{w}_{z_1(i)} - \mathbf{w}_{z_2(i)}\|_F^2 \geq \mu_0^2 |\{i : z_1(i) \neq z_2(i)\}| \frac{s}{2}.$$

Let $\epsilon_1^2 = (c_1 \log K)/48$. Since $\mu_0^2 = (c_1 \log K)/s$, we have $|\{i : z_1(i) \neq z_2(i)\}| \leq n/6$. Denote $B_\epsilon(\boldsymbol{\mu} \mathbf{L}^T) := \{\mathbf{L}_1 \in \mathcal{L}_K : d_1((\boldsymbol{\mu}, \mathbf{L}), (\boldsymbol{\mu}, \mathbf{L}_1)) \leq \epsilon\}$ for any $\boldsymbol{\mu} \in \mathbb{R}^{p \times K}$ and $\mathbf{L} \in \mathcal{L}_K$. Then for any $(\boldsymbol{\mu}, \mathbf{L}) \in \Theta_{K1}^*$, we have

$$\begin{aligned} |B_{\epsilon_1}(\boldsymbol{\mu} \mathbf{L}^T)| &\leq \binom{n}{n/6} K^{n/6} \leq \exp\left(n \log 6 - \frac{5n}{6} \log 5 + \frac{n}{6} \log K\right) \\ &= \exp\left(n \log 6 - \frac{5 \log 5}{6 \log 6} n \log 6 + \frac{n}{6} \log K\right) \leq \exp\left(\frac{3}{10} n \log 6 + \frac{n}{6} \log K\right) \\ &\leq \exp\left(\frac{3 \log 6}{10 \log 2} n \log K + \frac{1}{6} n \log K\right) \leq \exp(0.95n \log K), \end{aligned}$$

where the second inequality comes from the Stirling's formula and the fourth inequality is due to the fact that $K \geq 2$. Denote $M(\epsilon_1, \Theta_{K1}^*, d_1)$ as the ϵ -packing number of Θ_{K1}^* with respect to the metric d_1 . Since $|\Theta_{K1}^*| = K^n$, we have

$$\log M(\epsilon_1, \Theta_{K1}^*, d_1) \geq \log \frac{K^n}{\binom{n}{n/6} K^{n/6}} \geq \frac{1}{20} n \log K.$$

Note that $D_{KL}(\mathbb{P}_{\boldsymbol{\mu} \mathbf{L}_1^T} \| \mathbb{P}_{\boldsymbol{\mu} \mathbf{L}_2^T}) = \frac{1}{2} \|\boldsymbol{\mu} \mathbf{L}_1 - \boldsymbol{\mu} \mathbf{L}_2\|_F^2 \leq \frac{2\mu_0^2 sn}{2} = c_1 n \log K$. Therefore, by the generalized Fano's lemma (Yu, 1997)

$$\inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_{K1}^*} \mathbb{P}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \geq c_2 n \log K \right) \geq 1 - \frac{c_1 n \log K + \log 2}{n \log K / 20} \geq 0.9$$

with some constant $c_2 = c_1/48 > 0$ for sufficiently small $c_1 < 1/20$ and sufficiently large n where \mathbb{P}_* represents the probability measure under $(\boldsymbol{\mu}^*, \mathbf{L}^*, \boldsymbol{\Sigma}^*)$. By Markov's inequality,

$$\begin{aligned} & \inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_{K_1}^*} \mathbb{E}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \\ & \geq \inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_{K_1}^*} c_2 n \log K \mathbb{P}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \geq c_2 n \log K \right) \geq 0.9 c_2 n \log K. \end{aligned}$$

■ We next consider the minimax lower bound over $\Theta_{K_2}^*$. By the construction of $\Theta_{K_2}^*$, we have $\boldsymbol{\mu} \mathbf{L}^T = (\boldsymbol{\mu}_1 \ \dots \ \boldsymbol{\mu}_1 \ \boldsymbol{\mu}_2 \ \dots \ \boldsymbol{\mu}_2 \ \dots \ \boldsymbol{\mu}_K \ \dots \ \boldsymbol{\mu}_K)$. Thus, $\|\boldsymbol{\mu} \mathbf{L}^T - \boldsymbol{\mu}' \mathbf{L}'^T\|_F^2 = (n/K) \sum_{k=1}^K \|\boldsymbol{\mu}_k - \boldsymbol{\mu}'_k\|_2^2$. Let $\boldsymbol{\mu}_k = (\lambda_k \ a_0 \mathbf{v}^T)^T$ for $k \in [K]$, where λ_k 's are distinct scalars for $k \in [K]$, $\mathbf{v} \in \{0, 1\}^{p-1}$, and $a_0 > 0$ is to be specified later. By Lemma 4.10 in Massart (2007), there exists $\{\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(N)}\} \subset \{0, 1\}^{p-1}$ such that

- $\|\mathbf{v}^{(i)} - \mathbf{v}^{(j)}\|^2 > \frac{s-1}{2}$ for $i \neq j$,
- $\|\mathbf{v}^{(i)}\|_0 = s - 1$ for all $i \in [N]$,
- $N > \exp(c_4(s-1) \log \frac{p-1}{s-1})$ for some $c_4 \geq 0.233$.

For each distinct pair $\boldsymbol{\mu}, \boldsymbol{\mu}'$, we choose

$$\boldsymbol{\mu} = \begin{pmatrix} \lambda_1 & \dots & \lambda_K \\ a_0 \mathbf{v} & \dots & a_0 \mathbf{v} \end{pmatrix} \quad \text{and} \quad \boldsymbol{\mu}' = \begin{pmatrix} \lambda_1 & \dots & \lambda_K \\ a_0 \mathbf{v}' & \dots & a_0 \mathbf{v}' \end{pmatrix}$$

such that $\mathbf{v}, \mathbf{v}' \in \{\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(N)}\}$, $\mathbf{v} \neq \mathbf{v}'$. Then we consider an ϵ_2 -ball in $\Theta_{K_2}^*$ with respect to the metric $d_1((\boldsymbol{\mu}, \mathbf{L}_1), (\boldsymbol{\mu}', \mathbf{L}_2)) = \|\boldsymbol{\mu} \mathbf{L}_1^T - \boldsymbol{\mu}' \mathbf{L}_2^T\| / \sqrt{n}$. We have

$$\frac{1}{n} \|\boldsymbol{\mu} \mathbf{L}^T - \boldsymbol{\mu}' \mathbf{L}'^T\|_F^2 = \frac{1}{K} \sum_{k=1}^K \|\boldsymbol{\mu}_k - \boldsymbol{\mu}'_k\|_2^2 = \frac{1}{K} \sum_{k=1}^K a_0^2 \|\mathbf{v} - \mathbf{v}'\|_2^2 > \frac{1}{K} a_0^2 \frac{(s-1)K}{2} = \frac{a_0^2(s-1)}{2}.$$

Let $a_0^2 = 2\{sc_3 \log(p/s)\} / \{n(s-1)\}$ and $\epsilon_2^2 = \{c_3 s \log(p/s)\} / n$, where $c_3 > 0$ is a constant to be specified later. Then we have

$$M(\epsilon_2, \Theta_{K_2}^*, d_1) \geq \exp \left(c_4(s-1) \log \frac{p-1}{s-1} \right) \geq \exp \left(\frac{c_4}{4} s \log \frac{p}{s} \right)$$

for $s \geq 2$ because $(p-1)/(s-1) \geq \sqrt{p/s}$. Note that for any $(\boldsymbol{\mu}, \mathbf{L}), (\boldsymbol{\mu}', \mathbf{L}') \in \Theta_{K_2}^*$,

$$D_{KL}(\mathbb{P}_{\boldsymbol{\mu} \mathbf{L}^T} \| \mathbb{P}_{\boldsymbol{\mu}' \mathbf{L}'^T}) = \frac{1}{2} \|\boldsymbol{\mu} \mathbf{L} - \boldsymbol{\mu}' \mathbf{L}'^T\|_F^2 \leq a_0^2 s n.$$

Without loss of generality we assume $s \geq 11$. Then $\log 2 \leq \frac{10}{11} \frac{c_4}{4} s \log \left(\frac{p}{s} \right)$ since $c_4 \geq 0.2$, and

$$\frac{d_{KL}(\Theta_{K_2}^*) + \log 2}{\log M(\epsilon_2, \Theta_{K_2}^*, d_1)} \leq \frac{a_0^2 s n + \log 2}{\frac{c_4}{4} s \log \frac{p}{s}} \leq \frac{16c_3}{c_4} + \frac{10}{11}.$$

Therefore, by selecting $c_3 = c_4/(22 \times 16)$, we have

$$\inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_{K_2}^*} \mathbb{P}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \geq c_3 s \log \frac{p}{s} \right) \geq \frac{1}{22}$$

for $c_4 \geq 0.233$, and for sufficiently large n . It follows from Markov's inequality that

$$\inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_{K_2}^*} \mathbb{E}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \geq \frac{c_3}{22} s \log \frac{p}{s}.$$

■ Now we combined the minimax lower bounds over $\Theta_{K_1}^*$ and $\Theta_{K_2}^*$:

$$\begin{aligned} & \inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_K^*} \mathbb{E}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \\ & \geq \inf_{\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}} \max_{j \in \{1, 2\}} \sup_{(\boldsymbol{\mu}^*, \mathbf{L}^*) \in \Theta_{K_j}^*} \mathbb{E}_* \left(\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \\ & \geq \max \left\{ 0.9c_2 n \log K, \frac{c_3}{22} s \log \frac{p}{s} \right\} \geq C \left(s \log \frac{p}{s} + n \log K \right). \end{aligned}$$

■

Proof of Theorem 3 Note that by a basic inequality, we have

$$\begin{aligned} \|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 & \geq \|\mathbf{Y} - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T\|_F^2 = \|(\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T) + ((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T)\|_F^2 \\ & = \|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 + \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T\|_F^2 + 2\langle \mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T, (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T \rangle_F \\ & = \|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 + \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T\|_F^2 \\ & \quad + 2 \left\langle \mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T, \frac{(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T}{\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T\|_F} \right\rangle_F \times \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T\|_F, \end{aligned}$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product defined by $\langle \mathbf{A}, \mathbf{B} \rangle_F = \text{tr}(\mathbf{A}^T \mathbf{B})$. After rearranging the terms on the both sides of the inequality, we obtain

$$\mathbb{E}_* \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T\|_F \leq 2 \mathbb{E}_* \left[\sup_{(\boldsymbol{\mu}, \mathbf{L}) \in \Theta_K} \left\langle \mathbf{E}, (\boldsymbol{\Sigma}^*)^{\frac{1}{2}} \frac{\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T}{\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F} \right\rangle_F \right],$$

where $\mathbf{E} = (\boldsymbol{\Sigma}^*)^{-\frac{1}{2}}(\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)$ is the standardized noise matrix.

Consider the set of matrices

$$\tilde{\Theta}_K = \left\{ (\boldsymbol{\Sigma}^*)^{\frac{1}{2}} \frac{\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T}{\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F} : \boldsymbol{\mu} \in \mathbb{R}^{p \times K}, |\text{supp}(\boldsymbol{\mu})| \leq s, \mathbf{L} \in \mathcal{L}_K \right\}.$$

To obtain an upper bound of the right hand side of the inequality above, we use some tools of maximal inequality of empirical process. Specifically, we define a stochastic process $X(\mathbf{B}) = \langle \mathbf{E}, \mathbf{B} \rangle_F$ indexed by a $p \times n$ matrix \mathbf{B} . Since the entries of \mathbf{E} are i.i.d. standard Gaussian, it follows that $X(\mathbf{B})$ is sub-Gaussian. Then by Corollary 8.5 in Kosorok (2008),

$$\mathbb{E}_* \left[\sup_{(\boldsymbol{\mu}, \mathbf{L}) \in \Theta_K} \left\langle \mathbf{E}, (\boldsymbol{\Sigma}^*)^{\frac{1}{2}} \frac{\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T}{\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F} \right\rangle_F \right] \lesssim \int_0^{\text{diam}(\tilde{\Theta}_K)} \sqrt{\log N(\epsilon, \tilde{\Theta}_K, \|\cdot\|_F)} d\epsilon.$$

Obtaining a sharp upper bound of $N(\epsilon, \tilde{\Theta}_K, \|\cdot\|_F)$ is quite involved. We breakdown the computation of a sharp bound for the covering number of $\tilde{\Theta}_K$ as follows.

Step 1: Decompose $\tilde{\Theta}_K$ into unions of subspaces \mathcal{E}_j^K where $\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F$ is bounded. Define a function $B : \mathbb{R}^{p \times K} \times \mathcal{L}_K \mapsto \tilde{\Theta}_K$ as $B(\mu, \mathbf{L}) = (\Sigma^*)^{\frac{1}{2}} \frac{\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F}$. Then we have

$$\begin{aligned} \|B(\mu, \mathbf{L}) - B(\mu', \mathbf{L}')\|_F &\leq \left\| (\Sigma^*)^{\frac{1}{2}} \left\| \frac{\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F} - \frac{(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T}{\|(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T\|_F} \right\|_F \right\|_F \\ &= \lambda_{\max}(\Sigma^*)^{\frac{1}{2}} \left\| \frac{\mu\mathbf{L}^T \|(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T\|_F - (\mu')(\mathbf{L}')^T \|(\mu^*)(\mathbf{L}^*)^T - \mu\mathbf{L}^T\|_F}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F \|(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T\|_F} \right. \\ &\quad \left. - \frac{(\mu^*)(\mathbf{L}^*)^T (\|(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T\|_F - \|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F)}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F \|(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T\|_F} \right\|_F \\ &\leq \lambda_{\max}(\Sigma^*)^{\frac{1}{2}} \left(\frac{\|\mu\mathbf{L}^T - (\mu')(\mathbf{L}')^T\|_F}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F} \right. \\ &\quad \left. + \frac{\left| \|(\mu')(\mathbf{L}')^T - (\mu^*)(\mathbf{L}^*)^T\|_F - \|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F \right|}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F} \right) \\ &\leq 2\lambda_{\max}(\Sigma^*)^{\frac{1}{2}} \frac{\|\mu\mathbf{L}^T - (\mu')(\mathbf{L}')^T\|_F}{\|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F}. \end{aligned}$$

We split Θ_K as follows. Define

$$\mathcal{E}_j^K = \{(\mu, \mathbf{L}) : \mu \in \mathbb{R}^{p \times K}, |\text{supp}(\mu)| \leq s, \mathbf{L} \in \mathcal{L}_K, \|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F \in (a_j, a_{j+1}]\},$$

where the sequence $\{a_j\}_{j=-\infty}^{\infty}$ satisfies $\lim_{j \rightarrow \infty} a_j = \infty$ and $\lim_{j \rightarrow -\infty} a_j = 0$. We also require that $a_j \leq a_{j+1}/\sqrt{n}$. Suppose \mathcal{N}_j^K is an $a_j\epsilon/2$ -covering of \mathcal{E}_j^K with respect to the Frobenius norm. Then, it follows that $\bigcup_{j=-\infty}^{\infty} B(\mathcal{N}_j^K)$ is an $\epsilon\lambda_{\max}(\Sigma^*)^{\frac{1}{2}}$ -covering of $\tilde{\Theta}_K$ with respect to the Frobenius norm. Thus,

$$N(\epsilon\lambda_{\max}(\Sigma^*)^{\frac{1}{2}}, \tilde{\Theta}_K, \|\cdot\|_F) \leq \sum_{j=-\infty}^{\infty} N(a_j\epsilon/2, \mathcal{E}_j^K, \|\cdot\|_F).$$

Step 2: Show that \mathbf{L} equals \mathbf{L}^* up to a permutation for $(\mu, \mathbf{L}) \in \mathcal{E}_j^K$ when a_{j+1} is small. Denote Δ as the minimum distance among all cluster centers, that is, $\Delta = \min_{i,j \in [K]} \|\mu_i^* - \mu_j^*\|_2$. For a fixed \mathbf{L} which is induced by $z : [n] \mapsto [K]$, we denote $n_{gh} \stackrel{\Delta}{=} |z^{-1}(g) \cap (z^*)^{-1}(h)|$, $n_g = \sum_{h=1}^K n_{gh}$, and $n_h^* = \sum_{g=1}^K n_{gh}$. We then have

$$\begin{aligned} \|\mu\mathbf{L}^T - (\mu^*)(\mathbf{L}^*)^T\|_F^2 &= \sum_{g=1}^K \sum_{h=1}^K n_{gh} \|\mu_g - \mu_h^*\|_2^2 \\ &= \sum_{g=1}^K \left\{ n_g \left\| \mu_g - \frac{\sum_{h=1}^K n_{gh} \mu_h^*}{n_g} \right\|_2^2 + \sum_{h=1}^K n_{gh} \|\mu_h^*\|_2^2 - \frac{\left\| \sum_{h=1}^K n_{gh} \mu_h^* \right\|_2^2}{n_g} \right\} \\ &= \sum_{g=1}^K n_g \left\| \mu_g - \frac{\sum_{h=1}^K n_{gh} \mu_h^*}{n_g} \right\|_2^2 + C_{\mathbf{L}}, \end{aligned}$$

where $C_{\mathbf{L}} \triangleq \sum_{g=1}^K \left\{ \sum_{h=1}^K n_{gh} \|\boldsymbol{\mu}_h^*\|_2^2 - \left\| \sum_{h=1}^K n_{gh} \boldsymbol{\mu}_h^* \right\|_2^2 / n_g \right\}$.

Note that by Cauchy-Schwarz inequality and triangle inequality, we have

$$\left\| \sum_{h=1}^K n_{gh} \boldsymbol{\mu}_h^* \right\|_2^2 \leq \left(\sum_{h=1}^K n_{gh} \|\boldsymbol{\mu}_h^*\|_2 \right)^2 \leq \sum_{h=1}^K n_{gh} \sum_{h=1}^K n_{gh} \|\boldsymbol{\mu}_h^*\|_2^2 = n_g \sum_{h=1}^K n_{gh} \|\boldsymbol{\mu}_h^*\|_2^2$$

for every $g \in [K]$, which implies $C_{\mathbf{L}} \geq 0$ for any \mathbf{L} . Note that if $\mathbf{L}\mathbf{P} = \mathbf{L}^*$ for some permutation matrix $\mathbf{P} \in \mathcal{S}^{K \times K}$ where $\mathcal{S}^{K \times K}$ is the set of all $K \times K$ permutation matrices and τ which is the permutation function induced by \mathbf{P} , then for every $g \in [K]$, we have $n_g = n_{g, \tau(g)}$ and $n_{gh} = 0$ for $h \neq \tau(g)$. Thus $\left\| \sum_{h=1}^K n_{gh} \boldsymbol{\mu}_h^* \right\|_2^2 = n_{g, \tau(g)}^2 \|\boldsymbol{\mu}_{\tau(g)}^*\|_2^2 = n_g n_{g, \tau(g)} \|\boldsymbol{\mu}_{\tau(g)}^*\|_2^2 = n_g \sum_{h=1}^K n_{gh} \|\boldsymbol{\mu}_h^*\|_2^2$, which implies $C_{\mathbf{L}} = 0$. Next, by rearranging the terms in the expression of $C_{\mathbf{L}}$ we have

$$\begin{aligned} \sum_{h=1}^K n_g n_{gh} \|\boldsymbol{\mu}_h^*\|_2^2 - \left\| \sum_{h=1}^K n_{gh} \boldsymbol{\mu}_h^* \right\|_2^2 &= \sum_{h=1}^K n_{gh} (n_g - n_{gh}) \|\boldsymbol{\mu}_h^*\|_2^2 - \sum_{h_1=1}^K \sum_{h_2 \neq h_1}^K n_{gh_1} n_{gh_2} (\boldsymbol{\mu}_{h_1}^*)^T \boldsymbol{\mu}_{h_2}^* \\ &= \sum_{h_1=1}^K \sum_{h_2 \neq h_1}^K \{ n_{gh_1} n_{gh_2} \|\boldsymbol{\mu}_{h_1}^*\|_2^2 - n_{gh_1} n_{gh_2} (\boldsymbol{\mu}_{h_1}^*)^T \boldsymbol{\mu}_{h_2}^* \} = \sum_{h_1=1}^K \sum_{h_2 \neq h_1}^K \frac{1}{2} n_{gh_1} n_{gh_2} \|\boldsymbol{\mu}_{h_1}^* - \boldsymbol{\mu}_{h_2}^*\|_2^2. \end{aligned}$$

Suppose there is no permutation matrix $\mathbf{P} \in \mathcal{S}^{K \times K}$ such that $\mathbf{L}\mathbf{P} = \mathbf{L}^*$. Then there exists some $g' \in [K]$ such that for some $h'_1, h'_2 \in [K]$ and $h'_1 \neq h'_2$, $n_{g'h'_1} \geq 1$ and $n_{g'h'_2} \geq 1$. Furthermore, for such $g' \in [K]$ and h'_1, h'_2 , $n_{g'h'_1} n_{g'h'_2} \geq (n_{g'} - 1)/(K - 1)$. If this is not true, then we obtain $n_{g'h'_1} (n_{g'} - n_{g'h'_1}) = \sum_{h_2 \neq h'_1} n_{g'h'_1} n_{g'h_2} < n_{g'} - 1$ for the aforementioned h'_1 , and this contradicts to the fact that $n_{g'h'_1} (n_{g'} - n_{g'h'_1}) \geq n_{g'} - 1$ for $1 \leq n_{g'h'_1} \leq n_{g'} - 1$. The reason is that $n_{g'h'_1} (n_{g'} - n_{g'h'_1})$ is a quadratic function of $n_{g'h'_1}$ and the minimum is achieved when $n_{g'h'_1} = 1$ or $n_{g'h'_1} = n_{g'} - 1$. Therefore, we know that if \mathbf{L} is not identical to \mathbf{L}^* up to permutation, then there exists some $g' \in [K]$ and $h'_1 \neq h'_2$ such that

$$\begin{aligned} C_{\mathbf{L}} &= \sum_{g=1}^K \left(\sum_{h=1}^K n_{gh} \|\boldsymbol{\mu}_h^*\|_2^2 - \frac{\left\| \sum_{h=1}^K n_{gh} \boldsymbol{\mu}_h^* \right\|_2^2}{n_g} \right) = \sum_{g=1}^K \sum_{h_1=1}^K \sum_{h_2 \neq h_1}^K \frac{n_{gh_1} n_{gh_2} \|\boldsymbol{\mu}_{h_1}^* - \boldsymbol{\mu}_{h_2}^*\|_2^2}{2n_g} \\ &\geq \frac{n_{g'h'_1} n_{g'h'_2} \|\boldsymbol{\mu}_{h'_1}^* - \boldsymbol{\mu}_{h'_2}^*\|_2^2}{n_{g'}} \geq \frac{n_{g'} - 1}{n_{g'}(K - 1)} \Delta^2 \geq \frac{\Delta^2}{2(K - 1)}. \end{aligned}$$

Thus for \mathcal{E}_j^K , we can see that if $a_{j+1} \leq \Delta/\sqrt{2(K-1)}$, then every $(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E}_j^K$ satisfies $\mathbf{L}\mathbf{P} = \mathbf{L}^*$ for some permutation matrix $\mathbf{P} \in \mathcal{S}^{K \times K}$.

Step 3: Reduction of covering numbers of \mathcal{E}_j^K 's for small a_{j+1} . For $(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E}_j^K$ with $a_{j+1} \leq \Delta/\sqrt{2(K-1)}$, we have $\|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F = \|(\boldsymbol{\mu}\mathbf{P} - \boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F \leq \sqrt{n} \|\boldsymbol{\mu}\mathbf{P} - \boldsymbol{\mu}^*\|_F$ for some permutation matrix $\mathbf{P} \in \mathcal{S}^{K \times K}$. For a fixed $\mathbf{P} \in \mathcal{S}^{K \times K}$, denote $\mathcal{G}_j^K(\mathbf{P}) \triangleq \{\boldsymbol{\mu} \in \mathbb{R}^{p \times K} : |\text{supp}(\boldsymbol{\mu})| \leq s, \|\boldsymbol{\mu}\mathbf{P} - \boldsymbol{\mu}^*\|_F \in (a_j, a_{j+1}/\sqrt{n}]\}$, $\mathcal{G}_j^K \triangleq \{\boldsymbol{\mu} \in \mathbb{R}^{p \times K} : |\text{supp}(\boldsymbol{\mu})| \leq 2s, \|\boldsymbol{\mu}\|_F \in (a_j, a_{j+1}/\sqrt{n}]\}$. Then for every $\mathbf{P} \in \mathcal{S}^{K \times K}$, there is an injective

function $f_{\mathbf{P}} : \mathcal{G}_j^K(\mathbf{P}) \rightarrow \mathcal{G}_j^K$ such that $f_{\mathbf{P}}(\boldsymbol{\mu}) = \boldsymbol{\mu}\mathbf{P} - \boldsymbol{\mu}^*$. Thus f is a bijective function between $\mathcal{G}_j^K(\mathbf{P})$ and $R_f(\mathcal{G}_j^K(\mathbf{P}))$ where $R_f(\mathcal{G}_j^K(\mathbf{P})) \subset \mathcal{G}_j^K$ is the image of function f . Note that f is also distance-preserving with respect to the Frobenius norm, i.e., $\|f(\boldsymbol{\mu}) - f(\boldsymbol{\mu}')\|_F = \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_F$. Thus for any $\eta > 0$, $N(\eta, \mathcal{G}_j^K(\mathbf{P}), \|\cdot\|_F) = N(\eta, R_f(\mathcal{G}_j^K(\mathbf{P})), \|\cdot\|_F) \leq N(\eta, \mathcal{G}_j^K, \|\cdot\|_F)$. We know that for every $(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E}_j^K$ such that $a_{j+1} \leq \Delta/\sqrt{2(K-1)}$, there exists $\mathbf{P} \in \mathcal{S}^{K \times K}$ such that $\mathbf{L}\mathbf{P} = \mathbf{L}^*$. Suppose $\mathcal{N}_j^K(\mathbf{P})$ is a ζ/\sqrt{n} -covering of $\mathcal{G}_j^K(\mathbf{P})$. Then there exists $\tilde{\boldsymbol{\mu}} \in \mathcal{N}_j^K(\mathbf{P})$ such that $\|\tilde{\boldsymbol{\mu}}\mathbf{L} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F = \|(\tilde{\boldsymbol{\mu}}\mathbf{P} - \boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F \in (a_j, a_{j+1}]$, i.e., $(\tilde{\boldsymbol{\mu}}, \mathbf{L}) \in \mathcal{E}_j^K$, and $\|\boldsymbol{\mu}\mathbf{L} - \tilde{\boldsymbol{\mu}}\mathbf{L}\|_F \leq \sqrt{n}\|\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}\|_F \leq \zeta$, which means $\bigcup_{\mathbf{P} \in \mathcal{S}^{K \times K}} \mathcal{N}_j^K(\mathbf{P})$ is a ζ -covering of \mathcal{E}_j^K . Then we have

$$N(a_j\epsilon/2, \mathcal{E}_j^K, \|\cdot\|_F) \leq \sum_{\mathbf{P} \in \mathcal{S}^{K \times K}} N(a_j\epsilon/(2\sqrt{n}), \mathcal{G}_j^K(\mathbf{P}), \|\cdot\|_F) \leq K!N(a_j\epsilon/(2\sqrt{n}), \mathcal{G}_j^K, \|\cdot\|_F)$$

when $a_{j+1} \leq \Delta/\sqrt{2(K-1)}$. Note that for the covering number of the space \mathcal{E}_j^K , we abuse the notation by using $\|\cdot\|_F$ to denote the metric $d((\boldsymbol{\mu}_1, \mathbf{L}_1), (\boldsymbol{\mu}_2, \mathbf{L}_2)) = \|\boldsymbol{\mu}_1\mathbf{L}_1^T - \boldsymbol{\mu}_2\mathbf{L}_2^T\|_F$.

Step 4: Reduction of covering numbers of \mathcal{E}_j^K 's for large a_{j+1} . Next we consider the case when a_j is relatively large. Specifically, when $a_j \geq 2\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F$ and $(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E}_j^K$, we know that $\|\boldsymbol{\mu}\mathbf{L}^T\|_F \in (a_j - \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F, a_{j+1} + \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F) \subset (a_j/2, 2a_{j+1}]$. For \mathcal{E}_j^K , we can write it as $\mathcal{E}_j^K = \bigcup_{m=0}^{K-1} \mathcal{E}_{j,m}^K$ where $\mathcal{E}_{j,m}^K$ is the subset of \mathcal{E}_j^K whose \mathbf{L} is induced by a clustering with m empty clusters. Then for $\mathcal{E}_{j,m}^K$, it suffices to consider $\mathcal{F}_{j,m}^K \triangleq \{(\boldsymbol{\mu}_{-m}, \mathbf{L}_{-m}) : (\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E}_{j,m}^K\}$ where $\boldsymbol{\mu}_{-m}$ and \mathbf{L}_{-m} are the sub-matrices of $\boldsymbol{\mu}$ and \mathbf{L} by deleting the columns that correspond to the empty clusters respectively. For those j 's with $a_j \geq 2\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F$, we further have

$$\begin{aligned} \mathcal{F}_{j,m}^K &\subset \{(\boldsymbol{\mu}_{-m}, \mathbf{L}_{-m}) : \|\boldsymbol{\mu}_{-m}\mathbf{L}_{-m}^T\|_F \in (\frac{a_j}{2}, 2a_{j+1}], |\text{supp}(\boldsymbol{\mu}_{-m})| \leq s\} \\ &\subset \bigcup_{\mathbf{L}_{-m} \in \mathcal{L}_{K-m}} \mathcal{H}_{j,m}^K \times \{\mathbf{L}_{-m}\} \end{aligned}$$

because the singular values of \mathbf{L}_{-m} are between 1 and \sqrt{n} , where

$$\mathcal{H}_{j,m}^K \triangleq \{\boldsymbol{\mu}_{-m} \in \mathbb{R}^{p \times (K-m)} : |\text{supp}(\boldsymbol{\mu}_{-m})| \leq s, \|\boldsymbol{\mu}_{-m}\|_F \in (a_j/2, 2a_{j+1}/\sqrt{n})\}.$$

Since for any $(\boldsymbol{\mu}_{-m}, \mathbf{L}_{-m})$ and $(\boldsymbol{\mu}'_{-m}, \mathbf{L}_{-m}) \in \mathcal{H}_{j,m}^K \times \{\mathbf{L}_{-m}\}$ we have

$$\|\boldsymbol{\mu}_{-m}\mathbf{L}_{-m}^T - \boldsymbol{\mu}'_{-m}\mathbf{L}_{-m}^T\|_F \leq \|\mathbf{L}_{-m}\|_2 \|\boldsymbol{\mu}_{-m} - \boldsymbol{\mu}'_{-m}\|_F \leq \sqrt{n}\|\boldsymbol{\mu}_{-m} - \boldsymbol{\mu}'_{-m}\|_F,$$

it follows that

$$N(a_j\epsilon/2, \mathcal{E}_j^K, \|\cdot\|_F) \leq \sum_{m=0}^{K-1} N(a_j\epsilon/2, \mathcal{F}_{j,m}^K, \|\cdot\|_F) \leq \sum_{m=0}^{K-1} |\mathcal{L}_{K-m}| N(a_j\epsilon/(2\sqrt{n}), \mathcal{H}_{j,m}^K, \|\cdot\|_F).$$

Step 5: Computing covering numbers of \mathcal{E}_j^K 's for small and large a_{j+1} . We denote $a_{-1} = \Delta/\sqrt{2(K-1)}$, $a_1 = \Delta/\sqrt{2(K-1)}$, $a_2 = 2\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F$, $\mathcal{E}_0^K = \mathcal{E}_{-1}^K = \emptyset$ and

$$\frac{a_{j+1}}{a_j} = \begin{cases} \frac{\sqrt{n}}{4} \left(1 + \frac{1}{j^2}\right)^{\frac{1}{sK}} & \text{for } j > 2 \\ \sqrt{n} \left(1 + \frac{1}{(-j-1)^2}\right)^{\frac{1}{2sK}} & \text{for } j < -1 \end{cases}.$$

Note that without loss of generality we may assume $a_2 > a_1$. We have $a_{j+1}/\sqrt{n} \geq a_j$ and

$$\begin{aligned} \lim_{j \rightarrow \infty} a_j &= \lim_{j \rightarrow \infty} \left(\frac{n}{16}\right)^{\frac{j-1}{2}} \left(\prod_{i=1}^{j-1} \left(1 + \frac{1}{i^2}\right)\right)^{\frac{1}{sK}} 2\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F = \infty, \\ \lim_{j \rightarrow \infty} a_{-j} &= \lim_{j \rightarrow \infty} n^{-\frac{j-1}{2}} \left(\prod_{i=1}^{j-1} \left(1 + \frac{1}{i^2}\right)\right)^{-\frac{1}{2sK}} \frac{\Delta}{\sqrt{2(K-1)}} = 0. \end{aligned}$$

For $j > 2$, we have

$$\begin{aligned} N\left(\frac{a_j \epsilon}{2\sqrt{n}}, \mathcal{H}_{j,m}^K, \|\cdot\|_F\right) &\leq \binom{p}{s} \left(\frac{3\sqrt{n}}{a_j \epsilon}\right)^{s(K-m)} \left\{ \left(\frac{4a_{j+1}}{\sqrt{n}}\right)^{s(K-m)} - a_j^{s(K-m)} \right\} \\ &\leq \binom{p}{s} \left(\frac{3\sqrt{n}}{\epsilon}\right)^{sK} \left\{ \left(\frac{4a_{j+1}}{a_j \sqrt{n}}\right)^{sK} - 1 \right\} \leq \binom{p}{s} \left(\frac{3\sqrt{n}}{\epsilon}\right)^{sK} \frac{1}{j^2}. \end{aligned}$$

Then

$$\sum_{j=2}^{\infty} N\left(\frac{a_j \epsilon}{2}, \mathcal{E}_j^K, \|\cdot\|_F\right) \leq \sum_{j=1}^{\infty} \sum_{m=0}^{K-1} |\mathcal{L}_{K-m}| N\left(\frac{a_j \epsilon}{2\sqrt{n}}, \mathcal{H}_{j,m}^K, \|\cdot\|_F\right) \leq K^{n+1} \binom{p}{s} \left(\frac{3\sqrt{n}}{\epsilon}\right)^{sK} \frac{\pi^2}{6}.$$

Similarly,

$$\sum_{j=-\infty}^{-1} N\left(\frac{a_j \epsilon}{2}, \mathcal{E}_j^K, \|\cdot\|_F\right) \leq \sum_{j=-\infty}^{-1} K! N\left(\frac{a_j \epsilon}{2\sqrt{n}}, \mathcal{G}_j^K, \|\cdot\|_F\right) \leq K! \binom{p}{2s} \left(\frac{6\sqrt{n}}{\epsilon}\right)^{2sK} \frac{\pi^2}{6}.$$

Step 6: Computing the covering number of \mathcal{E}_1^K . Denote $\mathcal{F}_{1,m}^K(\mathbf{L}) \triangleq \{(\boldsymbol{\mu}_{-m}, \mathbf{L}_{-m}) : (\boldsymbol{\mu}_{-m}, \mathbf{L}_{-m}) \in \mathcal{E}_{1,m}^K\}$ for fixed \mathbf{L} which induces m empty clusters. Then we have $\mathcal{F}_{1,m}^K = \bigcup_{\mathbf{L} \in \mathcal{L}_{K-m}} \mathcal{F}_{1,m}^K(\mathbf{L})$. By the previous derivation we have that $\|\boldsymbol{\mu}_{-m} \mathbf{L}_{-m}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq a_{j+1}^2$ is equivalent to $\sum_{g=1}^{K-m} n_g \left\| \boldsymbol{\mu}_g - \frac{\sum_{h=1}^K n_{gh} \boldsymbol{\mu}_h^*}{n_g} \right\|_2^2 + C_{\mathbf{L}} \leq a_{j+1}^2$.

Denote $(\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \setminus S}, (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \cap S} \in \mathbb{R}^p$ as the vectors which have the same values as $\boldsymbol{\mu}_h^*$ on coordinates $\text{supp}(\boldsymbol{\mu}^*) \setminus S$ and $\text{supp}(\boldsymbol{\mu}^*) \cap S$ respectively, and 0 elsewhere. Then, for S which is the support of $\boldsymbol{\mu}$, we have $\boldsymbol{\mu}_h^* = (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \setminus S} + (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \cap S}$ and $\langle \boldsymbol{\mu}_g, (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \setminus S} \rangle = \langle (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \cap S}, (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \setminus S} \rangle = 0$. Thus, for \mathcal{E}_1^K we have

$$\sum_{g=1}^{K-m} n_g \left\| \boldsymbol{\mu}_g - \frac{\sum_{h=1}^K n_{gh} (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \cap S}}{n_g} \right\|_2^2 + \sum_{g=1}^{K-m} n_g \left\| \frac{\sum_{h=1}^K n_{gh} (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \setminus S}}{n_g} \right\|_2^2 + C_{\mathbf{L}} \leq a_2^2.$$

We then denote $C'_{\mathbf{L}} \triangleq C_{\mathbf{L}} + \sum_{g=1}^{K-m} n_g \left\| \frac{\sum_{h=1}^K n_{gh} (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \setminus S}}{n_g} \right\|_2^2$.

Denote $\mathcal{F}_{1,m,S}^K(\mathbf{L}) \triangleq \{(\boldsymbol{\mu}_{-m}, \mathbf{L}_{-m}) \in \mathcal{F}_{1,m}^K(\mathbf{L}) : \text{supp}(\boldsymbol{\mu}) = S, |S| \leq s\}$ for fixed $S \subset [p]$. Then we have $\mathcal{F}_{1,m}^K(\mathbf{L}) = \bigcup_{S \subset [p], |S| \leq s} \mathcal{F}_{1,m,S}^K(\mathbf{L})$ and therefore

$$N(a_1 \epsilon / 2, \mathcal{E}_1^K, \|\cdot\|_F) \leq \sum_{m=0}^K \sum_{\mathbf{L} \in \mathcal{L}_{K-m}} \sum_{S \subset [p], |S| \leq s} N(a_1 \epsilon / 2, \mathcal{F}_{1,m,S}^K(\mathbf{L}), \|\cdot\|_F).$$

Let $\mathcal{I}_{1,m,S}^K(\mathbf{L}) \triangleq \left\{ (\boldsymbol{\mu}, \mathbf{L}_{-m}) : \sum_{g=1}^{K-m} n_g \|\boldsymbol{\mu}_g\|_2^2 + C'_L \leq a_2^2, \text{supp}(\boldsymbol{\mu}) = S, |S| \leq s \right\}$ for fixed \mathbf{L} which induces m empty clusters. Note that there is an injective function $f : \mathcal{F}_{1,m,S}^K(\mathbf{L}) \rightarrow \mathcal{I}_{1,m,S}^K(\mathbf{L})$ such that $f(\boldsymbol{\mu}_g) = \boldsymbol{\mu}_g - \sum_{h=1}^K n_{gh} (\boldsymbol{\mu}_h^*)^{\text{supp}(\boldsymbol{\mu}^*) \cap S} / n_g$ for $g \in [K-m]$ and we know that $\mathcal{I}_{1,m,S}^K(\mathbf{L})$ is contained in an $s(K-m)$ -dimensional ellipsoid with center 0 and length of semi-axes $\{(a_2^2 - C'_L)/n_g\}_{g=1}^{K-m}$. Thus the volume of $\mathcal{F}_{1,m,S}^K(\mathbf{L})$ can be bounded

$$|\mathcal{F}_{1,m,S}^K(\mathbf{L})| \leq |\mathcal{I}_{1,m,S}^K(\mathbf{L})| \leq \frac{\pi^{\frac{s(K-m)}{2}}}{\Gamma\left(\frac{s(K-m)}{2} + 1\right)} \prod_{g=1}^{K-m} \frac{(a_2^2 - C'_L)^{s/2}}{n_g^{s/2}} \leq \frac{\pi^{\frac{s(K-m)}{2}}}{\Gamma\left(\frac{s(K-m)}{2} + 1\right)} \frac{a_2^{s(K-m)}}{\prod_{g=1}^{K-m} n_g^{s/2}}$$

where Γ is the Euler's Gamma function.

Suppose $\mathcal{M}_{1,m,S}^K(\mathbf{L})$ is a maximal $a_1\epsilon/2$ -packing of $\mathcal{F}_{1,m,S}^K(\mathbf{L})$ for fixed \mathbf{L} and $S \subset [p]$. Then for every $(\tilde{\boldsymbol{\mu}}_{-m}, \mathbf{L}_{-m}) \in \mathcal{M}_{1,m,S}^K(\mathbf{L})$, consider $U_{1,m,S}((\tilde{\boldsymbol{\mu}}_{-m}, \mathbf{L}_{-m}), a_1\epsilon/4) \subset \mathcal{F}_{1,m,S}^K$. We have $\|\boldsymbol{\mu}\mathbf{L}^T - \tilde{\boldsymbol{\mu}}\mathbf{L}^T\|_F^2 = \sum_{g=1}^{K-m} n_g \|\boldsymbol{\mu}_g - \tilde{\boldsymbol{\mu}}_g\|_2^2 \leq a_1^2\epsilon^2/16$. Let $\mathcal{U}_{1,m,S}(\mathbf{L}) = \left\{ (\boldsymbol{\mu}, \mathbf{L}_{-m}) \in \mathbb{R}^{p \times (K-m)} \times \mathcal{L}_{K-m} : \sum_{g=1}^{K-m} 16n_g \|\boldsymbol{\mu}_g\|_2^2 / (a_1^2\epsilon^2) \leq 1, \text{supp}(\boldsymbol{\mu}) = S, |S| \leq s \right\}$ for fixed \mathbf{L} and $S \subset [p]$. Since $\boldsymbol{\mu}$ shares the same support as $\tilde{\boldsymbol{\mu}}_{-m}$ for $(\boldsymbol{\mu}, \mathbf{L}_{-m}) \in U_{1,m,S}((\tilde{\boldsymbol{\mu}}_{-m}, \mathbf{L}_{-m}), a_1\epsilon/4)$, there exists a bijective function $f' : U_{1,m,S}((\tilde{\boldsymbol{\mu}}, \mathbf{L}), a_1\epsilon/4) \rightarrow \mathcal{U}_{1,m,S}(\mathbf{L})$ such that $f'(\boldsymbol{\mu}_g) = \boldsymbol{\mu}_g - \tilde{\boldsymbol{\mu}}_g$ for $g \in [K-m]$. In addition, we know that $\mathcal{U}_{1,m,S}(\mathbf{L})$ is essentially an $s(K-m)$ -dimensional ellipsoid with center 0 and length of semi-axes $\{(a_1^2\epsilon^2)/16n_g\}_{g=1}^{K-m}$.

Therefore, the volume $|U_{1,m,S}((\tilde{\boldsymbol{\mu}}, \mathbf{L}), \frac{a_1\epsilon}{4})| = \frac{\pi^{\frac{s(K-m)}{2}}}{\Gamma\left(\frac{s(K-m)}{2} + 1\right)} \prod_{g=1}^{K-m} \left(\frac{a_1\epsilon}{4\sqrt{n_g}}\right)^s$. Note that the sets $U_{1,m,S}((\tilde{\boldsymbol{\mu}}_{-m}, \mathbf{L}_{-m}), \frac{a_1\epsilon}{4})$ are disjoint when $\tilde{\boldsymbol{\mu}}_{-m}$ varies since $\mathcal{M}_{1,m,S}^K(\mathbf{L})$ is a packing. Then we have,

$$\begin{aligned} N\left(\frac{a_1\epsilon}{2}, \mathcal{F}_{1,m,S}(\mathbf{L}), \|\cdot\|_F\right) &\leq |\mathcal{M}_{1,m,S}^K(\mathbf{L})| \leq \frac{|\mathcal{F}_{1,m}^K|}{|U_{1,m,S}((\tilde{\boldsymbol{\mu}}, \mathbf{L}), a_1\epsilon/4)|} \\ &\leq \frac{\pi^{\frac{s(K-m)}{2}}/\Gamma\left(\frac{s(K-m)}{2} + 1\right)}{\pi^{\frac{s(K-m)}{2}}/\Gamma\left(\frac{s(K-m)}{2} + 1\right)} \left(\frac{2a_2}{\epsilon a_1}\right)^{s(K-m)} \leq \left(\frac{2}{\epsilon}\right)^{sK} \left(\frac{2\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F}{\Delta}\right)^{sK}. \end{aligned}$$

Therefore,

$$\begin{aligned} N(a_1\epsilon/2, \mathcal{E}_1^K, \|\cdot\|_F) &\leq \sum_{m=0}^{K-1} \sum_{\mathbf{L} \in \mathcal{L}_{K-m}} \sum_{S \in [p], |S| \leq s} |\mathcal{M}_{1,m,S}^K(\mathbf{L})| \\ &\leq KK^n \binom{p}{s} \left(\frac{2}{\epsilon}\right)^{sK} \left(\frac{2\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F}{\Delta}\right)^{sK} \lesssim KK^n \binom{p}{s} \left(\frac{2}{\epsilon}\right)^{sK} (sn)^{sK/2}. \end{aligned}$$

Note that $\Delta \geq 1/n^q$ and $\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 = O(sn)$ by assumptions. Then we have

$$\log \left(\sum_{j=-\infty}^{\infty} N\left(\frac{a_j\epsilon}{2}, \mathcal{E}_j^K, \|\cdot\|_F\right) \right) \lesssim n \log K + s \log \frac{p}{s} + sK \log n + sK \log \frac{6}{\epsilon}.$$

Therefore, by Corollary 8.5 in Kosorok (2008), we have

$$\begin{aligned}
 \mathbb{E}_* \left[\sup_{(\hat{\boldsymbol{\mu}}, \hat{\mathbf{L}}) \in \Theta_K} \left\langle \mathbf{E}, (\boldsymbol{\Sigma}^*)^{\frac{1}{2}} \frac{\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T}{\|\hat{\boldsymbol{\mu}} \hat{\mathbf{L}}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F} \right\rangle_F \right] &\lesssim \int_0^{\text{diam}(\tilde{\Theta}_K)} \sqrt{\log N(\epsilon, \tilde{\Theta}_K, \|\cdot\|_F)} d\epsilon \\
 &\lesssim \int_0^2 \sqrt{\log K + s \log \frac{p}{s} + n \log K + \frac{sK}{2} \log n + sK \log \frac{6\lambda_{\max}(\boldsymbol{\Sigma}^*)^{\frac{1}{2}}}{\epsilon}} d\epsilon \\
 &\leq \sqrt{s \log \frac{p}{s} + n \log K + sK \log n + 2\sqrt{sK} \int_0^1 \sqrt{-\log u + \log(12\lambda_{\max}(\boldsymbol{\Sigma}^*)^{\frac{1}{2}})} du} \\
 &\lesssim \sqrt{s \log \frac{p}{s} + n \log K}.
 \end{aligned}$$

■

Appendix B. Proofs for Section 3

B.1 Proof architecture

We first sketch the proof of Theorem 4 by providing technical lemmas below.

Lemma 11 *Under the conditions of Theorem 4, we have*

$$\Pi\{\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 < s \log p\} \geq \exp\{-c(s \log p + n \log K^*)\}.$$

Note that the prior of $\boldsymbol{\mu}$ is absolutely continuous with respect to the Lebesgue measure, which implies $|\text{supp}(\boldsymbol{\mu})| = p$ with probability 1. However, we expect most rows of $\boldsymbol{\mu}$ come from the “spike” distribution *a priori*, which implies the “magnitude” of these rows is quite small with high prior probability. This motivates us to define a generalized notation of the support. Formally, for $\delta > 0$, we define $\text{supp}_\delta(\boldsymbol{\mu}) \triangleq \{j \in [p] : \|\boldsymbol{\mu}_{j*}\|_1 \leq \delta\}$ as the soft support of $\boldsymbol{\mu}$ with threshold δ , where $\boldsymbol{\mu}_{j*}$ represents the j th row of $\boldsymbol{\mu}$. Let $(\boldsymbol{\mu}_k)_{S_\delta} = (\mu_{jk} : j \in \text{supp}_\delta(\boldsymbol{\mu})) \in \mathbb{R}^{|\text{supp}_\delta(\boldsymbol{\mu})|}$ denotes the sub-vector of $\boldsymbol{\mu}_k$ whose coordinates are in $\text{supp}_\delta(\boldsymbol{\mu})$. It is conceivable that for small δ , the size of the soft support of $\boldsymbol{\mu}$ is small compared with p with high prior probability. This heuristics is formalized through the following lemma.

Lemma 12 *Given K , suppose $\boldsymbol{\mu} \in \mathbb{R}^{p \times K}$ follows the prior specification (4) and (8) with some hyperparameters $\kappa > 0$, $\alpha > 1$, $\lambda_0 \gg \lambda_1 > 0$ and let $\delta = (1 + \kappa) \log p / \lambda_0$. Assume $K \log \log p \leq \log p$. Then given K , we have, for $\boldsymbol{\mu} \in \mathbb{R}^{p \times K}$,*

$$\Pi \left(|\text{supp}_\delta(\boldsymbol{\mu})| \geq \beta \left(s + \frac{n \log K}{\log p} \right) \right) \leq \exp(-c(s \log p + n \log K))$$

for some constants $\beta, c > 0$.

Lemma 13 *Let $(\boldsymbol{\mu}', \mathbf{L}') \in \Theta = \bigcup_{K=1}^{K_{\max}} \mathbb{R}^{p \times K} \times \mathcal{L}_K$ be such that $(\boldsymbol{\mu}')(\mathbf{L}')^T \neq (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$ and consider*

$$\mathcal{E} = \{(\boldsymbol{\mu}, \mathbf{L}) : \|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F \leq \delta \|\boldsymbol{\mu}'(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F\},$$

for some sufficiently small constant δ such that $0 < \delta < \|\Sigma^*\|_2 / (2\|\Sigma^*\|_2 + 4)$. Assume the conditions of Theorem 4 hold. Let $p_0(\mathbf{Y}_i)$ and $p^*(\mathbf{Y}_i)$ be the density functions of $N((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i, \mathbf{I}_p)$ and $N((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i, \Sigma^*)$ respectively. Denote $p_0(\mathbf{Y}) = \prod_{i=1}^n p_0(\mathbf{Y}_i)$ and $p^*(\mathbf{Y}) = \prod_{i=1}^n p^*(\mathbf{Y}_i)$. Then there exists a test function ϕ_n such that

$$\begin{aligned} \mathbb{E}_* \phi_n &\leq \exp\{-c_1 \|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2\}, \\ \sup_{(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{E}} \mathbb{E}_{(\boldsymbol{\mu}, \mathbf{L}, \mathbf{I}_p)} \left(\frac{p^*(\mathbf{Y})}{p_0(\mathbf{Y})} (1 - \phi_n) \right) &\leq \exp\{-c_2 \|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2\} \end{aligned}$$

where $c_1, c_2 > 0$ are some positive constants that are independent of n .

Lemma 14 *Let*

$$\mathcal{F}_n = \bigcup_{K=1}^{K_{\max}} \left\{ \boldsymbol{\mu} \mathbf{L}^T : \boldsymbol{\mu} \in \mathbb{R}^{p \times K}, |\text{supp}_\delta(\boldsymbol{\mu})| \leq \beta \left(s + \frac{n \log K}{\log p} \right), \max_{k \in [K]} \|(\boldsymbol{\mu}_k)_{S_\delta}\|_\infty \leq a_n, \mathbf{L} \in \mathcal{L}_K \right\}$$

where β and δ are defined as in Lemma 12 and $a_n = (s \log p + n \log K_{\max}) n^\gamma$ for some constant $\gamma > 0$. Denote $N(\epsilon_n, \mathcal{F}_n, d)$ as the covering number of \mathcal{F}_n with respect to the metric $d(\mathbf{A}, \mathbf{B}) = \|\mathbf{A} - \mathbf{B}\|_F / \sqrt{n}$. Suppose $\epsilon_n^2 = (s \log p + n \log K_{\max}) / n$. Then we have, for some constant c ,

$$N(\epsilon_n, \mathcal{F}_n, d) \leq \exp(c n \epsilon_n^2).$$

Lemma 15 *Let δ be defined as in Lemma 12 and \mathcal{F}_n be defined as in Lemma 14. Assume the conditions of Theorem 4 hold. Then we have, for some constant c ,*

$$\Pi(\mathcal{F}_n^c) \leq \exp\{-c(s \log p + n \log K^*)\}.$$

B.2 Proofs of the auxiliary lemmas

In this subsection, we provide the detailed proofs of the lemmas appearing in Section B.1.

Proof of Lemma 11 The proof of this lemma is based on a modification of that of Lemma 3.1 in Xie et al. (2022). Denote $\epsilon^2 = (s \log p) / n$. First by conditioning on the event $\{(\boldsymbol{\mu}, \mathbf{L}) \in \Theta : \mathbf{L} = \mathbf{L}^*, K = K^*\}$, we have that

$$\Pi\{\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 < n \epsilon^2 \mid \mathbf{L} = \mathbf{L}^*, K = K^*\} \geq \Pi\left(\bigcap_{k=1}^{K^*} \{\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_k^*\|_2^2 \leq \epsilon^2\}\right),$$

where n_k^* is the number of observations assigned to the k th cluster according to the cluster assignment matrix \mathbf{L}^* . Now we focus on the prior distribution of $\boldsymbol{\mu}_k$. Denote S_0 the true sparsity of $\boldsymbol{\mu}^*$. Note that for each $k \in [K^*]$, $\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_k^*\|_2 \leq \|(\boldsymbol{\mu}_k)_{S_0} - (\boldsymbol{\mu}_k^*)_{S_0}\|_2 + \|(\boldsymbol{\mu}_k)_{S_0^c}\|_2$,

$$\mathbb{P}\left(\bigcap_{k=1}^{K^*} \{\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_k^*\|_2^2 < \epsilon^2\}\right) \geq \mathbb{P}\left(\bigcap_{k=1}^{K^*} \left\{ \|(\boldsymbol{\mu}_k)_{S_0} - (\boldsymbol{\mu}_k^*)_{S_0}\|_2 < \frac{\epsilon}{2} \right\} \cap \left\{ \|(\boldsymbol{\mu}_k)_{S_0^c}\|_2 < \frac{\epsilon}{2} \right\} \right).$$

Now we introduce the latent random variable $\xi_j \sim \text{Bernoulli}(\theta)$ such that $|\mu_{jk}| \mid \xi_j \sim (1 - \xi_j) \text{Exp}(\lambda_0) + \xi_j \text{Exp}(\lambda_1)$ independently for all $j \in [p]$ and $k \in [K^*]$, where μ_{jk} is the

j th coordinate of $\boldsymbol{\mu}_k$. Recall that a Laplace distribution can be represented as a scale-mixture of normals as $(\mu_{jk} \mid \phi_j, \xi_j) \sim N\left(0, \phi_j/\lambda_{\xi_j}^2\right)$ with $\phi_j \sim \text{Exp}(1/2)$. Define the event $\mathcal{A} = \bigcap_{j \in S_0} \{\xi_j = 1\} \bigcap_{j \in S_0^c} \{\xi_j = 0\} \bigcap_{j \in S_0} \{1 < \phi_j < 4\}$. Note that given \mathcal{A} , the entries of $\boldsymbol{\mu}$ are independent. Conditioning on \mathcal{A} , we have $|\mu_{jk}| \sim \text{Exp}(\lambda_0)$ for $j \in S_0^c$, which implies

$$\begin{aligned} \prod_{k=1}^{K^*} \Pi\left(\|(\boldsymbol{\mu}_k)_{S_0^c}\|_2 < \frac{\epsilon}{2} \mid \mathcal{A}\right) &\geq \prod_{k=1}^{K^*} \prod_{j \in S_0^c} \Pi\left(|\mu_{jk}| < \frac{\epsilon}{2\sqrt{p}} \mid \mathcal{A}\right) \\ &= \left\{1 - \exp\left(-\frac{\lambda_0 \epsilon}{2\sqrt{p}}\right)\right\}^{K^*(p-s)} \geq \left(1 - \frac{s}{p}\right)^{pK^*} \geq \exp\{-\log(2e)sK^*\}. \end{aligned}$$

Here we use the inequality $(1-x)^{1/x} \geq 1/(2e)$ for $x \in (0, 1/2)$ and the fact $\lambda_0 \geq 2 \log \frac{p}{s} \sqrt{\frac{np}{s \log p}} \implies \frac{\lambda_0}{2} \sqrt{\frac{s \log p}{np}} \geq \log \frac{p}{s}$. Next, by Anderson's lemma (see Lemma 1.4 in supporting document for Pati et al. (2014)), for each $k \in [K^*]$, conditioning on \mathcal{A} (which guarantees that $1 \leq \phi_j \leq 4$ for all $j \in S_0$), we have

$$\begin{aligned} \Pi\left(\|(\boldsymbol{\mu}_k)_{S_0} - (\boldsymbol{\mu}_k^*)_{S_0}\|_2 < \frac{\epsilon}{2} \mid \mathcal{A}\right) &\geq \exp\left(-\frac{1}{2} \sum_{j \in S_0} \frac{|\mu_{jk}^*|^2 \lambda_1^2}{\phi_j}\right) \Pi\left(\|(\boldsymbol{\mu}_k)_{S_0}\|_2 < \frac{\epsilon}{2} \mid \mathcal{A}\right) \\ &\geq \exp\left(-\frac{1}{2} \sum_{j \in S_0} \frac{|\mu_{jk}^*|^2 \lambda_1^2}{\phi_j}\right) \prod_{j \in S_0} \left\{2\Phi\left(\frac{\epsilon \lambda_1}{2\sqrt{s\phi_j}}\right) - 1\right\} \\ &\geq \exp\left(-\frac{1}{2} \lambda_1^2 \|(\boldsymbol{\mu}_k^*)_{S_0}\|_2^2\right) \left\{2\Phi\left(\frac{\epsilon \lambda_1}{4\sqrt{s}}\right) - 1\right\}^s \geq \exp\left\{-\frac{1}{2} \lambda_1^2 \|(\boldsymbol{\mu}_k^*)_{S_0}\|_2^2 - s \left(1 + \left|\log \frac{\epsilon \lambda_1}{4\sqrt{s}}\right|\right)\right\}, \end{aligned}$$

where we use $\log(2\Phi(x) - 1) \geq -1 - |\log x|$ for small $x > 0$ in the last inequality. Since

$$\begin{aligned} sK^* \left|\log \frac{\epsilon \lambda_1}{4\sqrt{s}}\right| &= sK^* \left|\log \left(\frac{\lambda_1}{4} \sqrt{\frac{\log p}{n}}\right)\right| \\ &\leq sK^* \left|\log \frac{\lambda_1}{4}\right| + \frac{sK^*}{2} \left|\log \frac{\log p}{n}\right| \leq c'_1 sK^* \log n \leq c_1 s \log p \end{aligned}$$

for some constants $c'_1, c_1 > 0$, it follows that $\prod_{k=1}^{K^*} \Pi\left(\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_k^*\|_2^2 < \epsilon \mid \mathcal{A}\right) \geq \exp(-c_2 s \log p)$ for some constant $c_2 \geq \max(c_1, C) > 0$ given $\sum_{k=1}^{K^*} \lambda_1^2 \|(\boldsymbol{\mu}_k^*)_{S_0}\|_2^2 \leq C s \log p$. We next consider the prior probability of the event \mathcal{A} . First note that $\prod_{j \in S_0} \Pi(1 < \phi_j < 4) \geq \exp(-c_3 s)$ for some constant $c_3 > 0$ by the definition of the exponential distribution. Since the prior of θ is Beta(1, β_θ) where $\beta_\theta = p^{1+\kappa} \log p$, for some constant $c_4 > 1 + \kappa > 0$,

$$\begin{aligned} \Pi\left(\bigcap_{i \in S_0} \{\xi_i = 1\} \bigcap_{i \in S_0^c} \{\xi_i = 0\}\right) &= \int_0^1 \theta^s (1-\theta)^{p-s} \Pi(d\theta) \\ &= \frac{\Gamma(\beta_\theta + 1)}{\Gamma(\beta_\theta)} \int_0^1 \theta^s (1-\theta)^{p+\beta_\theta-s} d\theta = \frac{\Gamma(s+1)\Gamma(p+\beta_\theta-s)}{\Gamma(p+\beta_\theta+1)} \frac{\Gamma(\beta_\theta+1)}{\Gamma(\beta_\theta)} \\ &\geq \exp(-s \log(p+\beta_\theta) + \log \beta_\theta) \geq \exp(-s \log(2\beta_\theta)) \\ &\geq \exp(-s \log \log p - (1+\kappa)s \log p - s \log 2) \geq \exp(-c_4 s \log p). \end{aligned}$$

Hence, we obtain that $\Pi(\mathcal{A}) \geq \exp(-c_3s - c_4s \log p)$ and therefore,

$$\Pi \left(\bigcap_{k=1}^{K^*} \left\{ \|\boldsymbol{\mu}_k - \boldsymbol{\mu}_k^*\|_2^2 < \frac{\epsilon^2}{K^*} \right\} \right) \geq \prod_{k=1}^{K^*} \Pi \left(\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_k^*\|_2 < \frac{\epsilon}{\sqrt{K^*}} \mid \mathcal{A} \right) \Pi(\mathcal{A}) \geq \exp(-c_5s \log p)$$

for some constant $c_5 > \max(c_2, c_4) > 0$. Thus, we have

$$\Pi(\|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p \mid \mathbf{L} = \mathbf{L}^*, K = K^*) \geq \exp(-c_5s \log p).$$

Next consider $\Pi(\mathbf{L} = \mathbf{L}^* \mid K = K^*)$ for Multinomial-Dirichlet model. Let $\mathbf{w} = (w_1, \dots, w_{K^*})^T$. Define integers $\alpha_r = \lfloor \alpha \rfloor$ and $\beta_r = \lfloor K^* \alpha \rfloor$. Then we know $\lceil \alpha \rceil = \alpha_r + 1$ and $\lceil K^* \alpha \rceil = \beta_r + 1$. Note that the gamma function $\Gamma(x)$ is strictly increasing for $x > 2$ and we have $\Gamma(x) \leq 1$ for $1 \leq x \leq 2$. Thus, we have, for $K^* \alpha > 2$,

$$\frac{\Gamma(K^* \alpha)}{\Gamma(K^* \alpha + n)} \geq \frac{\Gamma(\beta_r)}{\Gamma(\beta_r + 1 + n)} = \frac{(\beta_r - 1)!}{(\beta_r + n)!};$$

for $1 \leq K^* \alpha \leq 2$,

$$\frac{\Gamma(K^* \alpha)}{\Gamma(K^* \alpha + n)} \geq \frac{\min_{x \in [1,2]} \Gamma(x)}{(\beta_r + n)!} = \left(\min_{x \in [1,2]} \Gamma(x) \right) \frac{\beta_r!}{(\beta_r + n)!}.$$

Similarly, we also have, for $\alpha > 2$,

$$\frac{\Gamma(\alpha + n_i)}{\Gamma(\alpha)} \geq \frac{\Gamma(\alpha_r + n_i)}{\Gamma(\alpha_r + 1)} \geq \frac{(\alpha_r + n_i - 1)!}{\alpha_r!};$$

for $1 \leq \alpha \leq 2$,

$$\frac{\Gamma(\alpha + n_i)}{\Gamma(\alpha)} \geq \frac{(\alpha_r + n_i - 1)!}{1} = \frac{(\alpha_r + n_i - 1)!}{\alpha_r!}.$$

Therefore, we conclude that for $\alpha \geq 1$ and some constants $C_1, C_2 > 0$,

$$\frac{\Gamma(K^* \alpha)}{\Gamma(K^* \alpha + n)} \geq C_1 \frac{(\beta_r - 1)!}{(\beta_r + n)!} \quad \text{and} \quad \frac{\Gamma(\alpha + n_i)}{\Gamma(\alpha)} \geq C_2 \frac{(\alpha_r + n_i - 1)!}{\alpha_r!}.$$

By integrating out \mathbf{w} we have

$$\begin{aligned} \Pi(\mathbf{L} = \mathbf{L}^* \mid K = K^*) &= \int \Pi(\mathbf{L} = \mathbf{L}^* \mid \mathbf{w}, K = K^*) d\Pi(\mathbf{w} \mid K = K^*) \\ &= \int_0^1 \cdots \int_0^1 \prod_{i=1}^n w_{z_i^*} \frac{\Gamma(K^* \alpha)}{(\Gamma(\alpha))^{K^*}} \prod_{j=1}^{K^*} w_j^{\alpha-1} dw_1 \cdots dw_{K^*} = \frac{\Gamma(K^* \alpha)}{(\Gamma(\alpha))^{K^*}} \frac{\prod_{i=1}^{K^*} \Gamma(\alpha + n_i)}{\Gamma(K^* \alpha + n)} \\ &\geq C_1 \frac{(\beta_r - 1)!}{(\beta_r + n)!} (C_2)^{K^*} \prod_{i=1}^{K^*} \frac{(\alpha_r + n_i - 1)!}{\alpha_r!} \\ &= C_1 (C_2)^{K^*} \frac{(\beta_r - 1)!}{(\beta_r + n)!} \frac{(K^* \alpha_r + n - K^*)!}{(\alpha_r!)^{K^*}} \frac{1}{\binom{K^* \alpha_r + n - K^*}{\alpha_r + n_1 - 1, \dots, \alpha_r + n_{K^*} - 1}} \end{aligned}$$

where $\binom{a}{a_1, \dots, a_{K^*}} = \frac{a!}{a_1! a_2! \dots a_{K^*}!}$ is the multinomial coefficient for $a_1 + a_2 + \dots + a_{K^*} = a$. We know that the function $(a_1, \dots, a_{K^*}) \mapsto \binom{a}{a_1, \dots, a_{K^*}}$ achieves the maximum when a_1, \dots, a_{K^*} are as close as possible to each other. Formally, let $a = qK^* + r$ where $q \triangleq a \bmod K^* = \lfloor a/K^* \rfloor$ and $0 \leq r < K^*$, then the multinomial coefficient is maximized when $a_1 = \dots = a_r = q + 1$ and $a_{r+1} = \dots = a_{K^*} = q$, and hence the maximal value is $a! / \{q!^{K^*-r} (q+1)!^r\}$. Then the preceding expression achieves the minimum when $n_1 = \dots = n_r = \lfloor n/K^* \rfloor + 1$ and $n_{r+1} = \dots = n_{K^*} = \lfloor n/K^* \rfloor$ where $r = n - \lfloor n/K^* \rfloor K^*$. Note that $\frac{\alpha_r}{\beta_r} = \frac{\lfloor \alpha \rfloor}{\lfloor K^* \alpha \rfloor} \geq \frac{\lfloor \alpha \rfloor}{(\lfloor \alpha \rfloor + 1)K^*} \geq \frac{1}{2K^*}$. So

$$\begin{aligned} \Pi(\mathbf{L} = \mathbf{L}^* \mid K = K^*) &\geq C_1(C_2)^{K^*} \frac{\{(\alpha_r + 1)(\alpha_r + 2) \cdots (\alpha_r + \lfloor \frac{n}{K^*} \rfloor - 1)\}^{K^*} (\lfloor \frac{n}{K^*} \rfloor + \alpha_r)^r}{\beta_r(\beta_r + 1)(\beta_r + 2) \cdots (\beta_r + n)} \\ &\geq C_1(C_2)^{K^*} \frac{(\alpha_r + 1)^{K^*} (\alpha_r + 2)^{K^*} \cdots (\alpha_r + \lfloor \frac{n}{K^*} \rfloor - 1)^{K^*}}{(\beta_r + 2K^*)^{K^*} (\beta_r + 4K^*)^{K^*} \cdots (\beta_r + 2\lfloor \frac{n}{K^*} \rfloor K^* - 2K^*)^{K^*}} \times \frac{1}{(\beta_r + n)^{r+K^*}} \\ &\geq C_1(C_2)^{K^*} \left(\frac{1}{2K^*}\right)^{\lfloor \frac{n}{K^*} \rfloor K^*} \times \frac{1}{(\beta_r + n)^{r+K^*}}. \end{aligned}$$

Also note the prior of K is a truncated Poisson, so we have

$$\Pi(K = K^*) \geq \frac{e^{-\lambda} \lambda^{K^*}}{K^*!} \geq \exp(-\lambda + K^* \log \lambda - K^* \log K^*) \geq \exp(-2K^* \log K^*).$$

Therefore, for some constant $c > 0$,

$$\begin{aligned} \Pi(\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p) &\geq \Pi(\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p \mid \mathbf{L} = \mathbf{L}^*, K = K^*) \Pi(\mathbf{L} = \mathbf{L}^* \mid K = K^*) \Pi(K = K^*) \\ &\geq \exp\{-c(s \log p + n \log K^*)\}. \end{aligned}$$

■

Proof proof of Lemma 12 We denote $\boldsymbol{\mu}_{j*}$ the j th row of $\boldsymbol{\mu}$ for $j = 1, \dots, p$. Note that from the prior model, we have $|\mu_{ji}| \mid \xi_j \sim (1 - \xi_j) \text{Exp}(\lambda_0) + \xi_j \text{Exp}(\lambda_1)$ where $\text{Exp}(\lambda)$ is the exponential distribution with parameter λ . Then we have $\|\boldsymbol{\mu}_{j*}\|_1 \mid \xi_j \sim (1 - \xi_j) \text{Gamma}(K, \lambda_0) + \xi_j \text{Gamma}(K, \lambda_1)$ where $\text{Gamma}(K, \lambda)$ is the Gamma distribution with shape K and rate λ . Thus, by the change of variable $u = \lambda_0 x$ and conditioning on the event $A = \{\theta \leq c_1(s + n \log K / \log p) / p^{1+\kappa}\}$ for some constant $c_1 > 0$, we have

$$\begin{aligned} \Pi(\|\boldsymbol{\mu}_{j*}\|_1 > \delta) &\leq (1 - \theta) \int_{\delta}^{\infty} \frac{\lambda_0^K}{\Gamma(K)} x^{K-1} e^{-\lambda_0 x} dx + \theta = (1 - \theta) \frac{\lambda_0^K}{\Gamma(K)} \frac{1}{\lambda_0^K} \int_{\lambda_0 \delta}^{\infty} u^{K-1} e^{-u} du + \theta \\ &< \frac{1}{\Gamma(K)} (\lambda_0 \delta)^K e^{-\lambda_0 \delta} + c_1 \frac{s + \frac{n \log K}{\log p}}{p^{1+\kappa}} < (\lambda_0 \delta)^K e^{-\lambda_0 \delta} + c_1 \frac{s + \frac{n \log K}{\log p}}{p^{1+\kappa}} \\ &\leq \exp(K \log(1 + \kappa) + K \log \log p - (1 + \kappa) \log p) + c_1 \frac{s + \frac{n \log K}{\log p}}{p^{1+\kappa}} \\ &\leq \exp(-\kappa \log p) + c_1 \frac{s + \frac{n \log K}{\log p}}{p^{1+\kappa}} \leq (1 + c_1) \frac{s + \frac{n \log K}{\log p}}{p^{1+\kappa}} \end{aligned}$$

for sufficiently large n . Note that the second inequality comes from the result of Natalini and Palumbo (2000) about upper incomplete gamma function for $\lambda_0 \delta = (1 + \kappa) \log p > K + 1$. The fourth inequality is due to the fact that the function $x \mapsto x^K e^{-x}$ is decreasing when $x > K$. Note that Hagerup and Rüb (1990) stated a version of Chernoff's inequality for binomial distributions that $\mathbb{P}(X > ap) \leq ((\frac{q}{a})^a \exp(a))^p$ if $X \sim \text{Binomial}(p, q)$ and $q \leq a < 1$. Then over the event $A = \{\theta \leq c_1(s + n \log K / \log p) / p^{1+\kappa}\}$ we have

$$\begin{aligned}
 & \Pi \left(|\text{supp}_\delta(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K}{\log p} \right) \mid A \right) \\
 & \leq \exp \left(-\beta \left(s + \frac{n \log K}{\log p} \right) \log \frac{\beta \left(s + \frac{n \log K}{\log p} \right)}{pq} + \beta \left(s + \frac{n \log K}{\log p} \right) \right) \\
 & \leq \exp \left(-\beta \left(s + \frac{n \log K}{\log p} \right) \log(\beta(1 + c_1)p^\kappa) \right) \\
 & \leq \exp \left(-\beta \left(\kappa s \log p + s \log(\beta(1 + c_1)) + \kappa n \log K + \frac{n \log K}{\log p} \log(\beta(1 + c_1)) \right) \right) \\
 & \leq \exp(-c_2(s \log p + n \log K))
 \end{aligned}$$

for some constant $0 < c_2 \leq \beta \kappa$. For the event A , we calculate the prior probability of A^c . Let $\beta_\theta = p^{1+\kappa} \log p$, we have

$$\begin{aligned}
 \Pi \left(\theta > c_1 \frac{s \log p + n \log K}{p^{1+\kappa} \log p} \right) &= \int_{c_1 \frac{s \log p + n \log K}{p^{1+\kappa} \log p}}^1 \frac{\Gamma(\beta_\theta + 1)}{\Gamma(\beta_\theta)} (1 - \theta)^{\beta_\theta - 1} d\theta \\
 &= \left(1 - c_1 \frac{s \log p + n \log K}{p^{1+\kappa} \log p} \right)^{\beta_\theta} \leq \exp \left(-\beta_\theta c_1 \frac{s \log p + n \log K}{p^{1+\kappa} \log p} \right) \\
 &\leq \exp(-c_1(s \log p + n \log K)).
 \end{aligned}$$

Therefore, for some constant $0 < c \leq \min(c_1, c_2)$ we have

$$\begin{aligned}
 \Pi \left(|\text{supp}_\delta(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K}{\log p} \right) \right) &= \int_0^1 \Pi \left(|\text{supp}_\delta(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K}{\log p} \right) \mid A \right) \Pi(d\theta) \\
 &\leq \int_0^{c_1 \frac{s \log p + n \log K}{p^{1+\kappa} \log p}} \Pi(|\text{supp}_\delta(\boldsymbol{\mu})| > \beta s \mid A) \Pi(d\theta) + \Pi \left(\theta > c_1 \frac{s \log p + n \log K}{p^{1+\kappa} \log p} \right) \\
 &\leq \exp(-c(s \log p + n \log K)).
 \end{aligned}$$

■

Proof of Lemma 13 Construct a likelihood ratio test $\phi_n = \mathbb{1}\{\|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 - \|\mathbf{Y} - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 \geq \rho\}$ where $\rho = (2m - 1) \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2$ for constant $m = \left(\|\boldsymbol{\Sigma}^*\|_2 - \sqrt{\|\boldsymbol{\Sigma}^*\|_2^2 - \|\boldsymbol{\Sigma}^*\|_2} \right) / 2 > 0$ if $\|\boldsymbol{\Sigma}^*\|_2 \geq 1$ and $m = \|\boldsymbol{\Sigma}^*\|_2 / 2$ if $\|\boldsymbol{\Sigma}^*\|_2 < 1$. For the type I error probability, we have under the true model, $\mathbf{Y}_i - ((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i \sim N(0, \boldsymbol{\Sigma}^*)$ for

each i . Thus, by Hoeffding's inequality for sub-Gaussian random variables,

$$\begin{aligned}
 & \mathbb{P}_* \left(\|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 - \|\mathbf{Y} - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 \geq \rho \right) \\
 &= \mathbb{P}_* \left(\sum_{i=1}^n (\mathbf{Y}_i - ((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i)^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i \geq m \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 \right) \\
 &\leq 2 \exp \left(-\frac{c_1 m^2 \|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^4}{\|(\boldsymbol{\Sigma}^*)^{\frac{1}{2}}((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)\|_F^2} \right) \\
 &\leq \exp(-c'_1 \|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2)
 \end{aligned}$$

for some constants $c_1, c'_1 > 0$ since $\frac{m^2}{\|\boldsymbol{\Sigma}^*\|_2} \geq \frac{1}{\|\boldsymbol{\Sigma}^*\|_2} \geq \frac{1}{M_{\boldsymbol{\Sigma}}} > 0$. For the type II error,

$$\begin{aligned}
 C &\triangleq \int \frac{p^*(\mathbf{Y})p(\mathbf{Y})}{p_0(\mathbf{Y})} d\mathbf{Y} = \mathbb{E}_* \frac{p(\mathbf{Y})}{p_0(\mathbf{Y})} = \mathbb{E}_* \exp \left(-\frac{1}{2} (\|\mathbf{Y} - \boldsymbol{\mu}\mathbf{L}^T\|_F^2 - \|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2) \right) \\
 &= \mathbb{E}_* \exp \left(\sum_{i=1}^n (\mathbf{Y}_i - ((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i)^T (\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i - \frac{1}{2} \|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \\
 &= \exp \left(\frac{1}{2} \|(\boldsymbol{\Sigma}^*)^{\frac{1}{2}}(\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)\|_F^2 - \frac{1}{2} \|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \\
 &\leq \exp \left(\frac{1}{2} (\|\boldsymbol{\Sigma}^*\|_2 - 1) \|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \\
 &\leq \exp \left(\frac{1}{2} (\|\boldsymbol{\Sigma}^*\|_2 - 1) (\|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F + \|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F)^2 \right) \\
 &\leq \exp \left(\frac{1}{2} (1 + \delta)^2 (\|\boldsymbol{\Sigma}^*\|_2 - 1) \|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) < \infty.
 \end{aligned}$$

Therefore,

$$\mathbb{E}_{(\boldsymbol{\mu}, \mathbf{L}, \mathbf{I}_p)} \left(\frac{p^*(\mathbf{Y})}{p_0(\mathbf{Y})} (1 - \phi_n) \right) = C \mathbb{P}_{\mathbf{Z}} \left(\|\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 - \|\mathbf{Z} - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 < \rho \right)$$

where \mathbf{Z} has density $\frac{p(\mathbf{Z})p^*(\mathbf{Z})}{C p_0(\mathbf{Z})}$. We then have

$$\begin{aligned}
 & \frac{p(\mathbf{Z})p^*(\mathbf{Z})}{p_0(\mathbf{Z})} \\
 &\propto \exp \left(-\frac{1}{2} \left(\|(\boldsymbol{\Sigma}^*)^{-\frac{1}{2}}(\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)\|_F^2 + \|\mathbf{Z} - \boldsymbol{\mu}\mathbf{L}^T\|_F^2 - \|\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \right) \right) \\
 &= \exp \left(-\frac{1}{2} \left(\|(\boldsymbol{\Sigma}^*)^{-\frac{1}{2}}(\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)\|_F^2 + \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T\|_F^2 \right. \right. \\
 &\quad \left. \left. + 2 \langle (\boldsymbol{\Sigma}^*)^{-\frac{1}{2}}(\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T), (\boldsymbol{\Sigma}^*)^{\frac{1}{2}}((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T) \rangle_F \right) \right) \\
 &\propto \exp \left(-\frac{1}{2} \left(\|(\boldsymbol{\Sigma}^*)^{-\frac{1}{2}}(\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T + \boldsymbol{\Sigma}^*((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T))\|_F^2 \right) \right).
 \end{aligned}$$

Thus $\mathbf{Z}_i \sim N(((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i - \boldsymbol{\Sigma}^*((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T)_i, \boldsymbol{\Sigma}^*)$ and $\mathbf{Z}_i - (\boldsymbol{\mu}\mathbf{L}^T)_i \sim N((\mathbf{I}_p - \boldsymbol{\Sigma}^*)((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T)_i, \boldsymbol{\Sigma}^*)$ for $i \in [n]$. Then, by the same argument as type I error,

$$\begin{aligned} & \mathbb{P}_{\mathbf{Z}}(\|\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 - \|\mathbf{Z} - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 < \rho) \\ &= \mathbb{P}_{\mathbf{Z}}\left(\sum_{i=1}^n (\mathbf{Z}_i - ((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i)^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i < m\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2\right) \\ &= \mathbb{P}_{\mathbf{Z}}\left(\sum_{i=1}^n (\mathbf{Z}_i - (\boldsymbol{\mu}\mathbf{L}^T)_i)^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i < m\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2\right. \\ &\quad \left.- \sum_{i=1}^n (\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i\right) \\ &\leq \mathbb{P}_{\mathbf{Z}}\left(\sum_{i=1}^n (\mathbf{Z}_i - (\boldsymbol{\mu}\mathbf{L}^T)_i)^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i < (\delta - 1 + m)\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2\right). \end{aligned}$$

Note that the inequality comes from Cauchy-Schwarz inequality,

$$\begin{aligned} & \langle \boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T, (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T \rangle_F \\ &= \langle \boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}')(\mathbf{L}')^T, (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T \rangle_F - \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 \\ &\leq (\delta - 1)\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2. \end{aligned}$$

Denote $T = \sum_{i=1}^n (\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i^T (\boldsymbol{\Sigma}^* - \mathbf{I}_p)((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i$ as the mean of random variable $\sum_{i=1}^n (\mathbf{Z}_i - (\boldsymbol{\mu}\mathbf{L}^T)_i)^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i$. Similarly we have

$$\begin{aligned} & -\langle \boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T, (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T \rangle_F \\ &= \langle \boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}')(\mathbf{L}')^T, (\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T \rangle_F + \|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 \\ &\leq (\delta + 1)\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2. \end{aligned}$$

Thus,

$$\begin{aligned} -T &= \sum_i (\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i^T (\boldsymbol{\Sigma}^* - \mathbf{I}_p)((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T)_i \\ &\leq \|\boldsymbol{\Sigma}^*\|_2 \langle (\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T), (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T \rangle_F \\ &\quad - \langle \boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T, (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T \rangle_F \\ &\leq (\|\boldsymbol{\Sigma}^*\|_2(\delta - 1) + (\delta + 1))\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2. \end{aligned}$$

Therefore, by Hoeffding's inequality,

$$\begin{aligned} & \mathbb{P}_{\mathbf{Z}}(\|\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 - \|\mathbf{Z} - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 < \rho) \\ &\leq \mathbb{P}_{\mathbf{Z}}\left(\sum_{i=1}^n (\mathbf{Z}_i - (\boldsymbol{\mu}\mathbf{L}^T)_i)^T ((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)_i - T\right. \\ &\quad \left.< (\delta - 1 + m + \|\boldsymbol{\Sigma}^*\|_2(\delta - 1) + \delta + 1)\|(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2\right) \\ &\leq 2 \exp\left(-\frac{c_2^2\|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^4}{2\|(\boldsymbol{\Sigma}^*)^{\frac{1}{2}}((\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)\|_F^2}\right) \leq 2 \exp(-c_2'\|(\boldsymbol{\mu}')(\mathbf{L}')^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2) \end{aligned}$$

for constant $c'_2 = c_2^2/(2\|\Sigma^*\|_2) > 0$ where we use the fact that $c_2 = \delta - 1 + m + \|\Sigma^*\|_2(\delta - 1) + \delta + 1 = (\|\Sigma^*\|_2 + 2)\delta - (\|\Sigma^*\|_2 - m) < 0$ since $\delta < \frac{\|\Sigma^*\|_2}{2(\|\Sigma^*\|_2 + 2)} < \frac{\|\Sigma^*\|_2 - m}{\|\Sigma^*\|_2 + 2}$. In addition, $\frac{c_2^2}{\|\Sigma^*\|_2} \geq \frac{(\|\Sigma^*\|_2 - m)^2}{\|\Sigma^*\|_2} \geq \frac{\|\Sigma^*\|_2}{4} \geq \frac{m_\Sigma}{4} > 0$. As a result,

$$\begin{aligned} \mathbb{E}_{(\boldsymbol{\mu}, \mathbf{L}, \mathbf{I}_p)} \left(\frac{p^*(\mathbf{Y})}{p_0(\mathbf{Y})} (1 - \phi_n) \right) &= C \mathbb{P}_{\mathbf{Z}} (\|\mathbf{Z} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 - \|\mathbf{Z} - (\boldsymbol{\mu}')(\mathbf{L}')^T\|_F^2 < \rho) \\ &\leq 2 \exp(-c_3 \|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2) \end{aligned}$$

for some constant

$$\begin{aligned} c_3 &= c'_2 - \frac{1}{2}(1 + \delta)^2 (\|\Sigma^*\|_2 - 1) \\ &= \frac{((\|\Sigma^*\|_2 + 2)\delta + m - \|\Sigma^*\|_2)^2 - \|\Sigma^*\|_2(\|\Sigma^*\|_2 - 1)(1 + \delta)^2}{2\|\Sigma^*\|_2} \\ &= \frac{(5\|\Sigma^*\|_2 + 4)\delta^2 + (-4\|\Sigma^*\|_2^2 + 2\|\Sigma^*\|_2 m + 4m - 2\|\Sigma^*\|_2)\delta + (m^2 - 2\|\Sigma^*\|_2 m + \|\Sigma^*\|_2)}{2\|\Sigma^*\|_2}. \end{aligned}$$

Denote $c_4(\delta) = 2\|\Sigma^*\|_2 c_3$ as a quadratic function of δ , then $c_4(0) = m^2 - 2\|\Sigma^*\|_2 m + \|\Sigma^*\|_2 > 0$. When $\|\Sigma^*\|_2 \geq 1$, $m = \frac{\|\Sigma^*\|_2 - \sqrt{\|\Sigma^*\|_2^2 - \|\Sigma^*\|_2}}{2} \leq \frac{1}{2}$. When $\|\Sigma^*\|_2 < 1$, $m = \frac{\|\Sigma^*\|_2}{2} \leq \frac{1}{2}$. Thus $c_4(0) \geq \frac{1}{4}$ and therefore there exists $\delta > 0$ such that $c_3 \geq \frac{1}{8\|\Sigma^*\|_2} \geq \frac{1}{8M_\Sigma} > 0$. \blacksquare

Proof of Lemma 14 Denote $\mathcal{F}_{nK} = \{(\boldsymbol{\mu}, \mathbf{L}) : \boldsymbol{\mu} \in \mathbb{R}^{p \times K}, (\boldsymbol{\mu}_k)_{S_\delta} \in [-a_n, a_n]^{\beta(s+n \log K/\log p)}$ for $k \in [K], |\text{supp}_\delta(\boldsymbol{\mu})| \leq \beta(s+n \log K/\log p), \mathbf{L} = \mathcal{L}_K, K \leq K_{\max}\}$. Then $N(\epsilon_n, \mathcal{F}_n, d) \leq \sum_{K=1}^{K_{\max}} N(\epsilon_n, \mathcal{F}_{nK}, d)$ since \mathcal{F}_{nK_1} and \mathcal{F}_{nK_2} are disjoint for $K_1 \neq K_2$. Consider for fixed \mathbf{L} ,

$$\frac{\|\boldsymbol{\mu} \mathbf{L}^T - \boldsymbol{\mu}' \mathbf{L}^T\|_F^2}{n} \leq \frac{\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_F^2 \|\mathbf{L}\|_F^2}{n} = \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_F^2.$$

Denote $\mathcal{G}_{nK1} \triangleq \{\boldsymbol{\mu} \in \mathbb{R}^{p \times K} : (\boldsymbol{\mu}_k)_{S_\delta} \in [-a_n, a_n]^{\beta(s+n \log K/\log p)}, k \in [K]\}$ and $\mathcal{G}_{nK2} \triangleq \mathcal{L}_K$. We know that the cardinality of \mathcal{G}_{nK2} is K^n . Therefore, $\log N(\epsilon_n, \mathcal{F}_{nK}, d) \leq n \log K + \log N(\epsilon_1, \mathcal{G}_{nK1}, \|\cdot\|_F)$ where $\epsilon_1^2 = (s \log p + n \log K_{\max})/n$. Let $\mathcal{G}_{nK1\delta^c} \triangleq \{\boldsymbol{\mu} \in \mathcal{G}_{nK1} : \boldsymbol{\mu}_{S_\delta} = 0\}$ and $\mathcal{G}_{nK1\delta} \triangleq \{\boldsymbol{\mu} \in \mathcal{G}_{nK1} : \boldsymbol{\mu}_{S_\delta^c} = 0\}$. Suppose \mathcal{N}_{δ^c} and \mathcal{N}_δ are the minimal $\epsilon_1/2$ -coverings of $\mathcal{G}_{nK1\delta^c}$ and $\mathcal{G}_{nK1\delta}$ respectively. Then for any $\boldsymbol{\mu} \in \mathcal{G}_{nK1}$, there exists $\tilde{\boldsymbol{\mu}} \in \mathcal{N}_{\delta^c}$ and $\bar{\boldsymbol{\mu}} \in \mathcal{N}_\delta$ such that $\|\boldsymbol{\mu} - (\tilde{\boldsymbol{\mu}} + \bar{\boldsymbol{\mu}})\|_F \leq \|\boldsymbol{\mu}_{S_\delta^c} - \tilde{\boldsymbol{\mu}}_{S_\delta^c}\|_F + \|\boldsymbol{\mu}_{S_\delta} - \bar{\boldsymbol{\mu}}_{S_\delta}\|_F \leq \epsilon_1$. Thus, we have $\log N(\epsilon_1, \mathcal{G}_{nK1}, \|\cdot\|_F) \leq \log N(\epsilon_1/2, \mathcal{G}_{nK1\delta^c}, \|\cdot\|_F) + \log N(\epsilon_1/2, \mathcal{G}_{nK1\delta}, \|\cdot\|_F)$.

Note that for $\mathcal{G}_{nK1\delta^c}$, we have $|\mathcal{G}_{nK1\delta^c}| < \delta^{pK}$. Since $\delta \lesssim \log p / (pK_{\max} \sqrt{n/\log p})$, we know that $|\mathcal{G}_{nK1\delta^c}| \lesssim |B_{\epsilon_1}^{p \times K}(0)|$ where $B_{\epsilon_1}^{p \times K}(0)$ is an ϵ_1 -ball in $\mathbb{R}^{p \times K}$. Thus $N(\epsilon_1/2, \mathcal{G}_{nK1\delta^c}, \|\cdot\|_F)$ is bounded above by some constant. We know that for a subset of Euclidean space,

$$\log N(\epsilon_1, \mathcal{G}_{nK1\delta}, \|\cdot\|_F) \leq \beta \left(s + \frac{n \log K}{\log p} \right) K \log \frac{3a_n}{\epsilon_1} + \log \left(\beta \left(s + \frac{n \log K}{\log p} \right) \right).$$

Note that since $(s + n \log K/\log p)/p \rightarrow 0$ as $n \rightarrow \infty$, by Stirling's formula we have $\log \left(\beta \left(s + \frac{n \log K}{\log p} \right) \right) \lesssim s \log p + n \log K$. By letting $a_n = (s \log p + n \log K_{\max}) n^\gamma$ we have

$\log N(\epsilon_1, \mathcal{G}_{nK1}, \|\cdot\|_F) \lesssim s \log p + n \log K$ since $K \log n \lesssim \log p$. Thus, $\log N(\epsilon_n, \mathcal{F}_{nK}, d) \lesssim s \log p + n \log K$. Therefore, for some constant $c' > 0$,

$$N(\epsilon_n, \mathcal{F}_n, d) \leq \sum_{K=1}^{K_{\max}} \exp(c(s \log p + n \log K)) \leq \exp(c'(s \log p + n \log K_{\max})) = \exp(c'n\epsilon_n^2).$$

■

Proof of Lemma 15 We have

$$\begin{aligned} \Pi(\mathcal{F}_n^c) &\leq \sum_{k=1}^{K_{\max}} \Pi \left(\bigcup_{i=1}^k \{\|\boldsymbol{\mu}_i\|_{\infty} > a_n\} \mid K = k \right) \Pi(K = k) \\ &\quad + \Pi \left(|\text{supp}_{\delta}(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K_{\max}}{\log p} \right) \right) \\ &\leq \sum_{k=1}^{K_{\max}} k p \Pi(|\mu_{11}| > a_n) \Pi(K = k) + \Pi \left(|\text{supp}_{\delta}(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K_{\max}}{\log p} \right) \right) \\ &\leq p \Pi(|\mu_{11}| > a_n) \sum_{k=1}^{K_{\max}} k \Pi(K = k) + \Pi \left(|\text{supp}_{\delta}(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K_{\max}}{\log p} \right) \right) \\ &\leq \frac{\lambda}{1 - e^{-\lambda}} p \Pi(|\mu_{11}| > a_n) + \Pi \left(|\text{supp}_{\delta}(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K_{\max}}{\log p} \right) \right). \end{aligned}$$

By lemma 12, the last term on the right hand side of the inequality is bounded above

$$\Pi \left(|\text{supp}_{\delta}(\boldsymbol{\mu})| > \beta \left(s + \frac{n \log K_{\max}}{\log p} \right) \right) \leq \exp(-c(s \log p + n \log K_{\max}))$$

for some constant $c > 0$. By the spike-and-slab lasso prior, we know that $|\mu_{11}| \sim (1 - \theta)\text{Exp}(\lambda_0) + \theta\text{Exp}(\lambda_1)$. Let $X = |\mu_{11}|$, we have $X \mid \theta \sim (1 - \theta)\text{Exp}(\lambda_0) + \theta\text{Exp}(\lambda_1)$. Then

$$\begin{aligned} \sup_{m \geq 1} \frac{(\mathbb{E}[|X|^m])^{1/m}}{m} &= \sup_{m \geq 1} \frac{(\mathbb{E}[\mathbb{E}[|X|^m \mid \theta]])^{1/m}}{m} = \sup_{m \geq 1} \frac{\mathbb{E} \left[(1 - \theta) \frac{m!}{\lambda_0^m} + \theta \frac{m!}{\lambda_1^m} \right]^{1/m}}{m} \\ &= \sup_{m \geq 1} \frac{1}{m} \left\{ \left(1 - \frac{1}{1 + \beta\theta} \right) \frac{m!}{\lambda_0^m} + \frac{1}{1 + \beta\theta} \frac{m!}{\lambda_1^m} \right\}^{1/m} \\ &\leq \sup_{m \geq 1} \frac{1}{m} \frac{(m!)^{1/m}}{\lambda_1} \leq \sup_{m \geq 1} \frac{1}{m} \frac{(em)^{1/m} m}{e\lambda_1} \leq \frac{2}{\lambda_1} \leq 2n^{\gamma} \leq \infty \end{aligned}$$

for any $n > 0$. Note that the first inequality is due to the power mean inequality for $\lambda_0 \geq \lambda_1$. Thus by Bernstein inequality, we have

$$\begin{aligned} \Pi(X \geq a_n) &= \Pi(X - \mathbb{E}X \geq a_n - \mathbb{E}X) \leq \Pi(X - \mathbb{E}X \geq \frac{a_n}{2}) \\ &\leq 2 \exp \left(-c \min \left(\frac{a_n^2}{4\|X\|_{\psi_1}}, \frac{a_n}{2\|X\|_{\psi_1}} \right) \right) \leq \exp(-c(s \log p + n \log K_{\max})). \end{aligned}$$

Note that here $\mathbb{E} X = \frac{1-\theta}{\lambda_0} + \frac{\theta}{\lambda_1} \leq \frac{a_n}{2}$ and $a_n = (s \log p + n \log K_{\max})n^\gamma$. ■

B.3 Proof of Theorem 4

Proof of Theorem 4 Denote $\bar{\epsilon}_n^2 = (s \log p + n \log K^*)/n$ and $\Theta = \bigcup_{K=1}^{K_{\max}} \mathbb{R}^{p \times K} \times \mathcal{L}_K$. Let d be a metric on Θ with $d((\boldsymbol{\mu}, \mathbf{L}), (\boldsymbol{\mu}', \mathbf{L}')) = \|\boldsymbol{\mu} \mathbf{L}^T - \boldsymbol{\mu}' \mathbf{L}'^T\|_F$. Let

$$U_n = \{(\boldsymbol{\mu}, \mathbf{L}) \in \Theta : d((\boldsymbol{\mu}, \mathbf{L}), (\boldsymbol{\mu}^*, \mathbf{L}^*)) < M\sqrt{n\bar{\epsilon}_n}\}.$$

By Bayes rule, we have $\Pi(U_n^c | \mathbf{Y}) = \frac{\int_{U_n^c} \frac{p_n(\mathbf{Y})}{p_0(\mathbf{Y})} d\Pi}{\int \frac{p_n(\mathbf{Y})}{p_0(\mathbf{Y})} d\Pi} := \frac{N_n}{D_n}$ where

$$p_0(\mathbf{Y}) = (2\pi)^{-\frac{np}{2}} \exp\left(-\frac{1}{2}\|\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2\right), p_n(\mathbf{Y}) = (2\pi)^{-\frac{np}{2}} \exp\left(-\frac{1}{2}\|\mathbf{Y} - \boldsymbol{\mu} \mathbf{L}^T\|_F^2\right).$$

We also denote $p^*(\mathbf{Y}) = (2\pi)^{-\frac{np}{2}} \det(\boldsymbol{\Sigma}^*)^{-\frac{n}{2}} \exp\left(-\frac{1}{2}\|(\boldsymbol{\Sigma}^*)^{-\frac{1}{2}}(\mathbf{Y} - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T)\|_F^2\right)$. By Lemma 11 we know that $\Pi(\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p) \geq \exp(-c_1 n \bar{\epsilon}_n^2)$ for some constant $c_1 > 0$. Denote $A_n = \{D_n > \exp(-c_2 n \bar{\epsilon}_n^2)\}$ for some $c_2 > c_1 > 0$. Thus,

$$A_n \supset \{D_n > \Pi(\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p) \exp(-(c_2 - c_1)n \bar{\epsilon}_n^2)\}.$$

Let $\mathbb{1}(A)$ denote the indicator random variable of event A . Then we can write

$$\begin{aligned} \mathbb{E}_* \Pi(U_n^c | \mathbf{Y}) &= \mathbb{E}_* \{(\phi_n + 1 - \phi_n) \Pi(U_n^c | \mathbf{Y})\} \\ &= \mathbb{E}_* \{\phi_n \Pi(U_n^c | \mathbf{Y})\} + \mathbb{E}_* \{(1 - \phi_n) \mathbb{1}(A_n) \Pi(U_n^c | \mathbf{Y})\} + \mathbb{E}_* \{(1 - \phi_n) \mathbb{1}(A_n^c) \Pi(U_n^c | \mathbf{Y})\} \\ &\leq \mathbb{E}_* \phi_n + \mathbb{E}_* \left\{ (1 - \phi_n) \exp(c_2 n \bar{\epsilon}_n^2) \int_{U_n^c} \frac{p_n(\mathbf{Y})}{p_0(\mathbf{Y})} d\Pi \right\} + \mathbb{P}_*(A_n^c). \end{aligned}$$

We treat the three terms on the right-hand side of the last equality separately. Denote

$$\mathcal{F}_n = \bigcup_{K=1}^{K_{\max}} \left\{ (\boldsymbol{\mu}, \mathbf{L}) : |\text{supp}_{\delta_\mu}(\boldsymbol{\mu})| \leq \beta \left(s + \frac{n \log K}{\log p} \right), \max_{k \in [K]} \|(\boldsymbol{\mu}_k)_{S_\delta}\|_\infty \leq a_n, \mathbf{L} \in \mathcal{L}_K \right\}$$

for δ_μ and β defined in Lemma 12, and $a_n = (s \log p + n \log K_{\max})n^\gamma$. Let $U_{n,j} = \{(\boldsymbol{\mu}, \mathbf{L}) \in \mathcal{F}_n : d((\boldsymbol{\mu}, \mathbf{L}), (\boldsymbol{\mu}^*, \mathbf{L}^*)) \in [j^2 n \bar{\epsilon}_n^2, (j+1)^2 n \bar{\epsilon}_n^2]\}$. Let $N_{n,j}$ be the maximal ϵ_n -nets $U_{n,j,1}, \dots, U_{n,j,N_j}$ that covers $U_{n,j}$ with respect to metric d . By Lemma 13 we have that for each $U_{n,j,h} \in N_{n,j}$, there exists a test $\phi_{n,j,h}$ such that $\mathbb{E}_* \phi_{n,j,h} \leq \exp(-c_3 n j^2 \bar{\epsilon}_n^2)$. Denote $\phi_n = \max_{j=M}^\infty \max_{h=1}^{|N_{n,j}|} \phi_{n,j,h}$. Then $\mathbb{E}_* \phi_n \leq \sum_{j=M}^\infty \sum_{h=1}^{|N_{n,j}|} \mathbb{E}_* \phi_{n,j,h}$. By Lemma 13, we have $\mathbb{E}_* \phi_n \leq N(\epsilon_n, \mathcal{F}_n, d) \frac{\exp(-c_3 n M^2 \bar{\epsilon}_n^2)}{1 - \exp(-c_3 n \bar{\epsilon}_n^2)}$. By lemma 14 we have $N(\epsilon_n, \mathcal{F}_n, d) \leq \exp(c_4 n \bar{\epsilon}_n^2)$. Thus the first term goes to 0 under \mathbb{E}_* as n tends to infinity and sufficiently large $M > 0$.

For the third term, by definition we have $\mathbb{P}_*(A_n^c) \leq 1 - \Pi\{D_n > \Pi(\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p) \exp(-(c_2 - c_1)n \bar{\epsilon}_n^2)\}$. Consider the event in the probability on the right hand side. By dividing both sides of the inequality by $\Pi(\|\boldsymbol{\mu} \mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p)$, we can rewrite

it in terms of Π' , which is the restricted and renormalized probability measure of prior Π conditioning on the event $E_n = \{\|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 \leq s \log p\}$. By Jensen's inequality,

$$\sum_i Z_{ni} \triangleq \int \sum_i \log \frac{p_n(\mathbf{Y}_i)}{p_0(\mathbf{Y}_i)} d\Pi' \leq \log \int \prod_i \frac{p_n(\mathbf{Y}_i)}{p_0(\mathbf{Y}_i)} d\Pi'.$$

Then we have that $\sum_i \mathbb{E}_* Z_{ni} = -\frac{1}{2} \int \|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F^2 d\Pi' \geq -\frac{1}{2} s \log p$. So for the event $\{D'_n > \exp(-(c_2 - c_1)n\bar{\epsilon}_n^2)\}$ where $D'_n = \int \prod_i \frac{p_n(\mathbf{Y}_i)}{p_0(\mathbf{Y}_i)} d\Pi'$, we have $\Pi(\log D'_n \geq -(c_2 - c_1)n\bar{\epsilon}_n^2) \geq \Pi(\sum_i Z_{ni} \geq -(c_2 - c_1)n\bar{\epsilon}_n^2) \geq \Pi(\sum_i (Z_{ni} - \mathbb{E}_* Z_{ni}) \geq -c'_2 n\bar{\epsilon}_n^2)$. Therefore, by Hoeffding inequality for sub-Gaussian random variable,

$$\begin{aligned} \Pi \left(\left| \sum_i Z_{ni} - \mathbb{E}_* Z_{ni} \right| \geq \frac{c'_2 n\bar{\epsilon}_n^2}{2} \right) &\leq 2 \exp \left(-c_5 \frac{n^2 \bar{\epsilon}_n^4}{\|(\boldsymbol{\Sigma}^*)^{\frac{1}{2}}((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T)\|_F^2} \right) \\ &\leq 2 \exp \left(-c_5 \frac{n^2 \bar{\epsilon}_n^4}{\|(\boldsymbol{\Sigma}^*)^{\frac{1}{2}}\|_2^2 \|((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T - \boldsymbol{\mu}\mathbf{L}^T)\|_F^2} \right) \leq \exp(-c'_5 n\bar{\epsilon}_n^2) \end{aligned}$$

for some constants $c_5, c'_5 > 0$ since $\|\boldsymbol{\Sigma}^*\|_2$ is upper bounded from infinity. Thus, as $n \rightarrow \infty$, $\mathbb{E}_* \mathbb{1}(A_n^c) \leq 1 - \Pi(\log D'_n \geq -(c_2 - c_1)n\bar{\epsilon}_n^2) \leq \exp(-c'_5 n\bar{\epsilon}_n^2) \rightarrow 0$.

For the second term, we have, by Fubini's theorem,

$$\mathbb{E}_*(1 - \phi_n) \exp(c_2 n\bar{\epsilon}_n^2) \int_{U_n^c} \frac{p_n(\mathbf{Y})}{p_0(\mathbf{Y})} d\Pi \leq \exp(c_2 n\bar{\epsilon}_n^2) \left(\int_{U_n^c \cap \mathcal{F}_n} \mathbb{E}_n \left((1 - \phi_n) \frac{p^*(\mathbf{Y})}{p_0(\mathbf{Y})} \right) d\Pi + \Pi(\mathcal{F}_n^c) \right).$$

By Lemma 15, we know that $\log \Pi(\mathcal{F}_n^c) \lesssim -n\bar{\epsilon}_n^2$. By Lemma 13, we have $\mathbb{E}_n \left((1 - \phi_n) \frac{p^*(\mathbf{Y})}{p_0(\mathbf{Y})} \right) \leq \exp(-c_5 n M^2 \bar{\epsilon}_n^2)$ so the above term goes to 0 as n tends to infinity. ■

B.4 Proof of Theorem 8

Proof of Theorem 8 We find the singular value decomposition (SVD) of $\boldsymbol{\mu}\mathbf{L}^T = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T$ for some diagonal matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{p \times n}$ and $\mathbf{U} \in \mathbb{O}^p$ and $\mathbf{V} \in \mathbb{O}^n$ where \mathbb{O}^m denotes the set of m by m orthogonal matrices. Consider the matrix \mathbf{L} , let $\boldsymbol{\Sigma}_L = \mathbf{L}^T \mathbf{L}$ be the diagonal matrix whose k th diagonal entry is the size of cluster k . Then denote $\mathbf{L}_N = \mathbf{L}\boldsymbol{\Sigma}_L^{-1/2}$ as the ‘‘normalization’’ of \mathbf{L} since it is orthogonal. On the other hand, for matrix $\boldsymbol{\mu}$, we suppose the corresponding QR decomposition is $\boldsymbol{\mu} = \mathbf{Q}\mathbf{R}$ for some $\mathbf{Q} \in \mathbb{O}^p$ and upper triangular matrix $\mathbf{R} \in \mathbb{R}^{p \times K}$. Then suppose the SVD of $\mathbf{R}\boldsymbol{\Sigma}_L^{1/2}$ is $\mathbf{R}\boldsymbol{\Sigma}_L^{1/2} = \mathbf{U}_1 \boldsymbol{\Sigma}_1 \mathbf{V}_1^T$ for some $\mathbf{U}_1 \in \mathbb{O}^p$ and $\mathbf{V}_1 \in \mathbb{O}^K$. Therefore we obtain $\boldsymbol{\mu}\mathbf{L}^T = \mathbf{Q}\mathbf{U}_1 \boldsymbol{\Sigma}_1 \mathbf{V}_1^T \mathbf{L}_N^T$ and we know that v_i , which denotes the i th column of $\mathbf{L}_N \mathbf{V}_1$, satisfies $\boldsymbol{\mu}\mathbf{L}^T v_i = \sigma_i u_i$ for $i = 1, \dots, K$ where u_i is the i th column of $\mathbf{Q}\mathbf{U}_1$ and σ_i is the i th singular value of $\boldsymbol{\mu}\mathbf{L}^T$.

Then we can use a variant of Davis-Kahan theorem (Yu et al., 2014). Suppose $\boldsymbol{\mu}\mathbf{L}^T = \mathbf{Q}\mathbf{U}_1 \boldsymbol{\Sigma}_1 \mathbf{V}_1^T \mathbf{L}_N^T$ and $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T = \mathbf{Q}^* \mathbf{U}_1^* \boldsymbol{\Sigma}_1^* (\mathbf{V}_1^*)^T (\mathbf{L}_N^*)^T$. Denote $D_F(\mathbf{O}_1, \mathbf{O}_2) = \inf_{\mathbf{V} \in \mathbb{O}^r} \|\mathbf{O}_1 - \mathbf{O}_2 \mathbf{V}\|_F$ for $\mathbf{O}_1, \mathbf{O}_2 \in \mathbb{O}^r$ and let $\|\sin \Theta(\mathbf{O}_1, \mathbf{O}_2)\|_F$ be the (Frobenius) sine-theta distance

between $\text{span}(\mathbf{O}_1)$ and $\text{span}(\mathbf{O}_2)$. Then the relationship between the metric D_F and sine-theta distance holds: $\|\sin \Theta(\mathbf{O}_1, \mathbf{O}_2)\|_F \leq D_F(\mathbf{O}_1, \mathbf{O}_2) \leq \sqrt{2}\|\sin \Theta(\mathbf{O}_1, \mathbf{O}_2)\|_F$. Note that for the right singular subspace of $\boldsymbol{\mu}\mathbf{L}^T$, we have

$$\begin{aligned} \sqrt{2}\|\sin \Theta(\mathbf{L}_N \mathbf{V}_1, \mathbf{L}_N^* \mathbf{V}_1^*)\|_F &\geq \inf_{\mathbf{V} \in \mathbb{O}^K} \|\mathbf{L}_N \mathbf{V}_1 - \mathbf{L}_N^* \mathbf{V}_1^* \mathbf{V}\|_F \\ &= \inf_{\mathbf{V} \in \mathbb{O}^K} \|(\mathbf{L}_N - \mathbf{L}_N^* \mathbf{V}_1^* \mathbf{V} (\mathbf{V}_1)^{-1}) \mathbf{V}_1\|_F \\ &= \inf_{\mathbf{V} \in \mathbb{O}^K} \|\mathbf{L}_N - \mathbf{L}_N^* \mathbf{V}_1^* \mathbf{V} (\mathbf{V}_1^*)^{-1}\|_F \geq \|\sin \Theta(\mathbf{L}_N, \mathbf{L}_N^*)\|_F. \end{aligned}$$

Then by Theorem 3 in Yu et al. (2014), we have

$$\|\sin \Theta(\mathbf{L}_N, \mathbf{L}_N^*)\|_F \leq \frac{2\sqrt{2}(2\sigma_{\max} + \|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_2)}{\sigma_{\min}^2} \times \|\boldsymbol{\mu}\mathbf{L}^T - (\boldsymbol{\mu}^*)(\mathbf{L}^*)^T\|_F$$

where σ_{\max} and σ_{\min} represent the max and min singular values of $(\boldsymbol{\mu}^*)(\mathbf{L}^*)^T$ respectively.

We denote $\mathbf{L}_N = [(l_N)_1, \dots, (l_N)_n]^T$ and $\mathbf{L}_N^* = [(l_N)_1^*, \dots, (l_N)_n^*]^T$. Note that \mathbf{L}_N^* and \mathbf{L}_N have at most K distinct rows. Let ζ be the minimum distance among these K distinct rows of \mathbf{L}_N^* with respect to ℓ_2 norm: $\zeta = \min_{(l_N)_i^* \neq (l_N)_j^*} \|(l_N)_i^* - (l_N)_j^*\|_2$. Let $\mathbf{O} = \arg \inf_{\mathbf{V} \in \mathbb{O}^K} \|\mathbf{L}_N - \mathbf{L}_N^* \mathbf{V}\|_F$. Define the set $\mathcal{I} = \{i : \|(l_N)_i - \mathbf{O}^T (l_N)_i^*\|_2 \geq \zeta/2\}$. Assume that the event $\mathcal{E}_n = \{\|\sin \Theta(\mathbf{L}_N, \mathbf{L}_N^*)\|_F \leq \eta_n\}$ occurs *a posteriori*, where $\eta_n = \frac{8\sqrt{2}\sigma_{\max}}{\sigma_{\min}^2} (M(s \log p + n \log K_{\max}))^{1/2}$ and M is the constant in Theorem 4. By Theorem 4, we know that $\mathbb{E}_*\{\Pi(\mathcal{E}_n)\} \rightarrow 1$ as $n \rightarrow \infty$. This implies that $|\mathcal{I}| \leq 4\eta_n^2/\zeta^2$ since otherwise we have $\|\mathbf{L}_N - \mathbf{L}_N^* \mathbf{O}\|_F^2 \geq (\zeta^2/4)(4\eta_n^2/\zeta^2) = \eta_n^2$, which contradicts with the definition of \mathcal{E}_n . Thus, for any $i, j \in \mathcal{I}^c$ with $(l_N)_i = (l_N)_j$, $\|(l_N)_i^* - (l_N)_j^*\|_2 \leq \|(l_N)_i - \mathbf{O}^T (l_N)_i^*\|_2 + \|(l_N)_j - \mathbf{O}^T (l_N)_j^*\|_2 < \zeta$, which implies $(l_N)_i^* = (l_N)_j^*$ since ζ is the minimum distance between pair of distinct rows of \mathbf{L}_N^* . On the other hand, note that $\zeta^2 = 1/n_{\max}^* + 1/n_2$ where n_2 is the second largest cluster size. Consequently, since

$$\begin{aligned} \sigma_{\min}((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T) &= \sqrt{\lambda_{K^*}((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T (\mathbf{L}^*) (\boldsymbol{\mu}^*)^T)} \geq \sqrt{n_{\min}^* \lambda_{K^*}((\boldsymbol{\mu}^*)(\boldsymbol{\mu}^*)^T)} \geq n_{\min}^* \sigma_{\min}(\boldsymbol{\mu}^*), \\ \sigma_{\max}((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T) &= \sqrt{\lambda_{\max}((\boldsymbol{\mu}^*)(\mathbf{L}^*)^T (\mathbf{L}^*) (\boldsymbol{\mu}^*)^T)} \leq \sqrt{n_{\max}^* \lambda_{\max}((\boldsymbol{\mu}^*)(\boldsymbol{\mu}^*)^T)} \leq n_{\max}^* \sigma_{\max}(\boldsymbol{\mu}^*), \end{aligned}$$

we have $|\mathcal{I}| \leq \frac{4\eta_n^2}{\zeta^2} \asymp \frac{(n_{\max}^*)^3 \sigma_{\max}(\boldsymbol{\mu}^*)^2}{(n_{\min}^*)^4 \sigma_{\min}(\boldsymbol{\mu}^*)^4} (s \log p + n \log K_{\max})$. Note that $n_k^* > |\mathcal{I}|$ for all k , namely, $\{(l_N)_i^* : i \notin \mathcal{I}\}$ consists of all K^* distinct rows of \mathbf{L}_N^* . Therefore, each of the unique $B_{\|\cdot\|_2}(\mathbf{O}^T (l_N)_i^*, \zeta/2)$ for $i \in [n]$, which is the ℓ_2 ball centered at $\mathbf{O}^T (l_N)_i^*$ with radius $\zeta/2$, contains at least one element of $\{(l_N)_i : i \in \mathcal{I}^c\}$. Recall that ζ is the minimum distance between any pair of distinct rows of \mathbf{L}_N^* , so these open balls are disjoint for distinct rows of \mathbf{L}_N^* . It follows from the pigeonhole principle that each ball contains exactly one element of $\{(l_N)_i : i \in \mathcal{I}^c\}$. Thus, if $(l_N)_i^* = (l_N)_j^*$ for $i, j \in \mathcal{I}^c$, then $(l_N)_i, (l_N)_j \in B_{\|\cdot\|_2}(\mathbf{O}^T (l_N)_i^*, \zeta/2)$, implying that $(l_N)_i = (l_N)_j$ by the fact that every such ball contains exactly one row of \mathbf{L} .

Therefore we prove that for any $i \in \mathcal{I}^c$, $(l_N)_i = (l_N)_j$ if and only if $(l_N)_i^* = (l_N)_j^*$. So the number of mis-clustered points are at most $|\mathcal{I}|$, which gives us the result because $\mathbb{E}_*\{\Pi(\mathcal{E}_n)\} \rightarrow 1$ as $n \rightarrow \infty$ by Theorem 4. \blacksquare

Appendix C. Posterior Inference via Gibbs Sampling

In this section, we introduce a Gibbs sampler for posterior inference of the proposed Bayesian sparse Gaussian mixture model. We design a sampler based on the algorithm proposed in Miller and Harrison (2018). Let \mathcal{C} denote the partition of $[n]$ according to the cluster memberships \mathbf{z} . Formally, $\mathcal{C} = \{E_k : |E_k| > 0\}$ where $E_k = \{i : z_i = k \text{ for } i \in [n]\}$. Let \mathcal{C}_{-i} be the partition of $[n] \setminus \{i\}$ according to the cluster memberships $\{z_j\}_{j \neq i}$. We also denote n_c as the number of data points in $c \in \mathcal{C}$ and n_c^- as the number of data points in $c \in \mathcal{C}_{-i}$. We can derive an urn representation for the mixture model from the exchangeable partition distribution: $\pi(\mathcal{C}) = V_n(|\mathcal{C}|) \prod_{c \in \mathcal{C}} \alpha^{(|c|)}$, where $V_n(|\mathcal{C}|) = \sum_{k=1}^{\infty} p_K(k) k_{(|\mathcal{C}|)} / (\alpha k)^{(n)}$ and p_K is the prior of K . Here $x^{(m)} \triangleq x(x+1) \cdots (x+m-1)$, $x_{(m)} \triangleq x(x-1) \cdots (x-m+1)$.

To address the non-conjugacy of the Laplace distribution, we re-write the $\text{SSL}(\lambda_0, \lambda_1, \theta)$ prior through the normal-scale-mixture representation of Laplace as follows: for $j \in [p]$,

$$(x_j \mid \phi_j, \xi_j = a) \sim N\left(0, \frac{\phi_j}{\lambda_a^2}\right), a = 0, 1, \quad \phi_j \sim \text{Exp}\left(\frac{1}{2}\right), \quad \xi_j \sim \text{Bernoulli}(\theta).$$

We obtain the following closed-form full conditional posterior distributions of $\boldsymbol{\mu}_c$, ϕ_c and ξ :

$$\begin{aligned} ((\boldsymbol{\mu}_c)_j \mid -) &\sim N\left(\sum_{l \in c} (\mathbf{Y}_l)_j \left(n_c + \frac{\lambda_{\xi_j}^2}{(\phi_c)_j}\right)^{-1}, \left(n_c + \frac{\lambda_{\xi_j}^2}{(\phi_c)_j}\right)^{-1}\right), \\ ((\phi_c)_j \mid -) &\sim \text{GiG}(0.5, (\boldsymbol{\mu}_c)_j^2 \lambda_{\xi_j}^2, 1), \\ (\xi_j \mid -) &\sim \text{Bernoulli}(\theta'), \quad \text{where } \theta' \propto \prod_{c \in \mathcal{C}} \frac{\lambda_1}{\sqrt{(\phi_c)_j}} \exp\left(-\frac{1}{2} \lambda_1^2 \frac{(\boldsymbol{\mu}_c)_j^2}{(\phi_c)_j}\right) \theta, \\ (\theta \mid -) &\sim \text{Beta}\left(1 + \sum_{j=1}^p \xi_j, \beta_\theta + p - \sum_{j=1}^p \xi_j\right). \end{aligned}$$

Here, $\text{GiG}(\zeta, \chi, \tau)$ denotes the generalized inverse Gaussian distribution whose probability density function is $f(x) = x^{\zeta-1} \exp(-(\chi/x + \tau x)/2)$. We also remark that there exists a potential label switching phenomenon when sampling centers $\boldsymbol{\mu}_c$ and auxiliary variables ϕ_c for all clusters. This can be prevented by the following alignment process.

- (i) Collect B post-burn-in samples $\boldsymbol{\mu}^{(b)} = \left(\boldsymbol{\mu}_1^{(b)} \quad \dots \quad \boldsymbol{\mu}_{K^{(b)}}^{(b)}\right)$ and $\mathbf{L}^{(b)} = \left(\mathbf{l}_1^{(b)} \quad \dots \quad \mathbf{l}_{K^{(b)}}^{(b)}\right)^T$ for $b = 1, \dots, B$, where $K^{(b)} \triangleq |\mathcal{C}^{(b)}|$ is the number of clusters in b th iteration.
- (ii) Find the index b^* that corresponds to the maximizer of the log-likelihood function: $b^* = \arg \min_{b \in [B]} \|\mathbf{Y} - (\boldsymbol{\mu}^{(b)})(\mathbf{L}^{(b)})^T\|_{\text{F}}^2$.
- (iii) For $b = 1, \dots, B$, find $\mathcal{P}^{(b)} = \arg \min_{\mathcal{P} \in \mathcal{S}^{K^{(b)} \times K^{(b)}}} \|\boldsymbol{\mu}^{(b^*)} - \boldsymbol{\mu}^{(b)} \mathcal{P}\|_{\text{F}}^2$, where $\mathcal{S}^{K^{(b)} \times K^{(b)}}$ is the set of all $K^{(b)} \times K^{(b)}$ permutation matrices.
- (iv) For $b = 1, \dots, B$, replace $\boldsymbol{\mu}^{(b)}$ by $\boldsymbol{\mu}^{(b)} \mathcal{P}^{(b)}$ and $\mathbf{L}^{(b)}$ by $\mathbf{L}^{(b)} \mathcal{P}^{(b)}$.

We provide the detailed Gibbs sampler in Algorithm 1 below. The R code can be found at <https://github.com/YanxunXu/HighDimClustering>.

Algorithm 1 The Gibbs sampler

Require: Initialization of $\mathcal{C}, \{\boldsymbol{\mu}_c : c \in \mathcal{C}\}, \boldsymbol{\xi}, \{\phi_c : c \in \mathcal{C}\}$

```

1: for  $b \leftarrow 1$  to  $B$  do
2:   for  $i \leftarrow 1$  to  $n$  do
3:      $t \leftarrow |\mathcal{C}_{-i}|$ 
4:     if  $z_i \neq z_l$  for all  $l \neq i$  then
5:       Remove  $\boldsymbol{\mu}_{z_j}$ 
6:     end if
7:     Sample  $\phi_{t+1} \sim p_\phi(\phi_{t+1})$ 
8:     Sample  $\boldsymbol{\mu}_{t+1} \sim p_{\boldsymbol{\mu}|\boldsymbol{\xi},\phi}(\boldsymbol{\mu}_{t+1} \mid \boldsymbol{\xi}, \phi_{t+1})$ 
9:     for  $k \leftarrow 1$  to  $t$  do
10:       $m_k \leftarrow (n_k^- + \alpha)p(\mathbf{Y}_i \mid \boldsymbol{\mu}_c)$  where  $n_k^-$  is the size of cluster  $k$  in  $\mathcal{C}_{-j}$ 
11:    end for
12:     $V_n(t) \leftarrow \frac{t!}{n!} \frac{\Gamma(\alpha t)}{n^{\alpha t - 1}} p_K(t)$ 
13:     $V_n(t+1) \leftarrow \frac{(t+1)!}{n!} \frac{\Gamma(\alpha(t+1))}{n^{\alpha(t+1) - 1}} p_K(t+1)$ 
14:     $m_{t+1} \leftarrow \alpha \frac{V_n(t+1)}{V_n(t)} p(\mathbf{Y}_i \mid \boldsymbol{\mu}_{t+1})$ 
15:    Sample  $z_i \sim \text{Categorical} \left( \frac{m_1}{\sum_{k=1}^{t+1} m_k}, \dots, \frac{m_{t+1}}{\sum_{k=1}^{t+1} m_k} \right)$ 
16:  end for
17:  for  $c \leftarrow 1$  to  $|\mathcal{C}|$  do
18:    for  $j \leftarrow 1$  to  $p$  do
19:      Sample  $(\boldsymbol{\mu}_c)_j \sim N \left( \sum_{l \in \mathcal{C}} (\mathbf{Y}l)_j \left( n_c + \frac{\lambda_{\xi_j}^2}{(\phi_c)_j} \right)^{-1}, \left( n_c + \frac{\lambda_{\xi_j}^2}{(\phi_c)_j} \right)^{-1} \right)$ 
20:    end for
21:  end for
22:  for  $c \leftarrow 1$  to  $|\mathcal{C}|$  do
23:    for  $j \leftarrow 1$  to  $p$  do
24:      Sample  $(\phi_c)_j \sim \text{GiG}(0.5, (\boldsymbol{\mu}_c)_j^2 \lambda_{\xi_j}^2, 1)$ 
25:    end for
26:  end for
27:  for  $j \leftarrow 1$  to  $p$  do
28:    
$$\theta' \leftarrow \frac{\prod_{c \in \mathcal{C}} \lambda_1 \exp \left( -\frac{1}{2} \lambda_1^2 \frac{(\boldsymbol{\mu}_c)_j^2}{(\phi_c)_j} \right) \theta}{\prod_{c \in \mathcal{C}} \lambda_1 \exp \left( -\frac{1}{2} \lambda_1^2 \frac{(\boldsymbol{\mu}_c)_j^2}{(\phi_c)_j} \right) \theta + \prod_{c \in \mathcal{C}} \lambda_0 \exp \left( -\frac{1}{2} \lambda_0^2 \frac{(\boldsymbol{\mu}_c)_j^2}{(\phi_c)_j} \right) (1 - \theta)}$$

29:    Sample  $\xi_j \sim \text{Bernoulli}(\theta')$ 
30:  end for
31:  Sample  $\theta \sim \text{Beta} \left( 1 + \sum_{j=1}^p \xi_j, \beta_\theta + p - \sum_{j=1}^p \xi_j \right)$ 
32: end for

```

Appendix D. Additional Numerical Studies

We conduct three additional simulation studies to evaluate the proposed method under various setups: (a) when the covariance matrix of the true sampling distribution is diagonal, with different diagonal entries; (b) when the true sampling distribution exhibits skewness; and (c) when the true coordinates do not exhibit sparsity. We use the same hyperparameters as in the simulation studies, and run the proposed method with 1000 burn-in samples and 4000 post-burn-in samples. Each additional simulation is replicated 100 times.

In (a), we modify Scenario I in the simulation study with $s = 12$ and $K = 3$ by setting $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$, where σ_j^2 are i.i.d. sampled from $\text{Gamma}(100, 100)$ for $j \in [p]$. In all 100 replicated simulations, the proposed method successfully identifies the three clusters, yielding an average ARI of 0.97. In contrast, the four competitors tend to merge the two overlapping clusters, resulting in average ARIs of 0.54. These results highlight the robustness of the proposed Bayesian method, consistently outperforming alternative methods in terms of clustering accuracy.

In (b), we consider the sampling distribution to be a skewed t-distribution, with the same location and scale parameters as Scenario III in the manuscript. We set the skewness parameter to 10 and the degrees of freedom to 2. The average ARI across these 100 replicates is 0.84 for our proposed method, compared to 0.34 for PCA-KM. MClust and SKM fail to cluster in almost all simulation replicates, returning an average ARI of less than 0.01.

In (c), we set $n = 200$, $p = 400$, $K = 2$, $(\mu_1^*)_S = r\mathbf{1}_s$, $(\mu_2^*)_S = -r\mathbf{1}_s$, and assume $\Sigma_1^* = \Sigma_2^* = \mathbf{I}_p$. We fix the “signal-to-noise ratio” by setting $\|\mu_1^* - \mu_2^*\|_2^2 = 36$ and the sparsity support is defined as $S = [s]$, with s varying over $\{4, 10, 100, 400\}$. Our proposed Bayesian model successfully recovers the two clusters when $s = 4$ and 10. However, when $s = 100$ and 400, our model returns only one cluster, indicating poorer performance when the underlying truth is denser. Intuitively, this discrepancy arises because the spike-and-slab prior inherently introduces bias when the true model is not sparse. This bias affects the estimation of cluster centers, thereby influencing the update of cluster membership during subsequent MCMC iterations.

	Scenario I	
	$K = 3, s = 6$	$K = 3, s = 12$
Bayesian	369	368
PCA-KM	0.04	0.03
MClust	1.27	1.51
SKM	22.61	21.73
CHIME	11.28	11.37

Table A1: Empirical running times (seconds) for all methods in simulation studies.

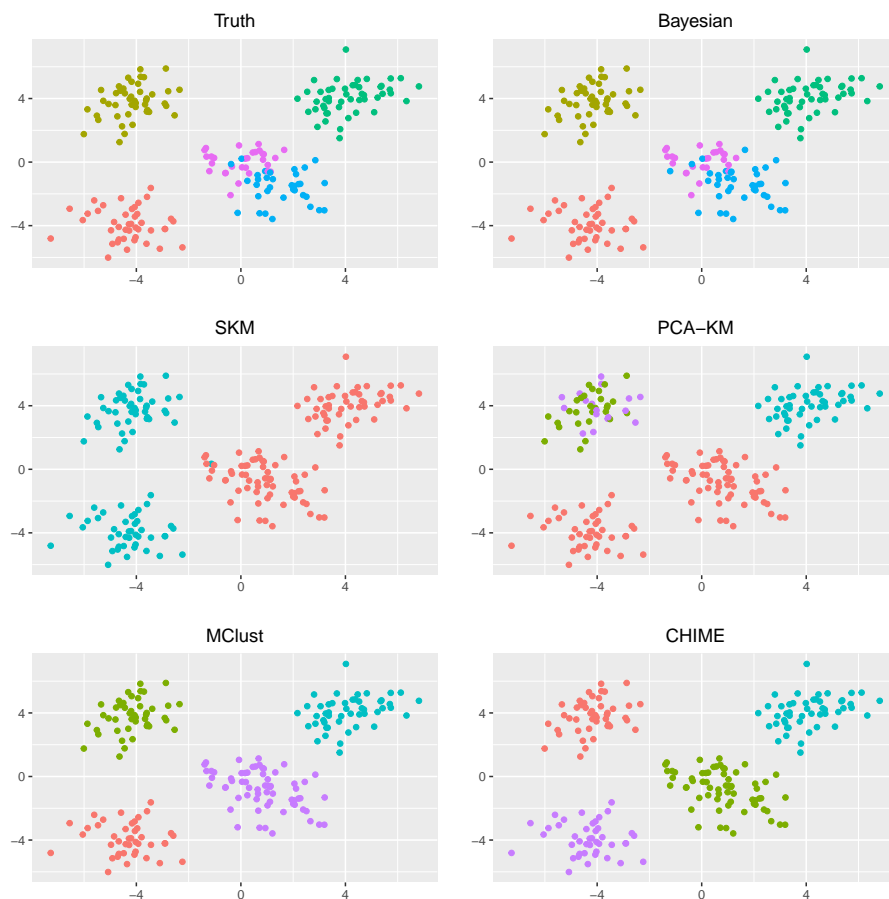


Figure A1: Clustering results of different methods compared with the true cluster memberships in the simulation Scenario I with $K^* = 5$ and $s = 6$ in a randomly selected simulation replicate. Data points are projected onto the subspace of the first two coordinates and different colors correspond to different estimated cluster memberships of the data points.

Methods	ARI	NMI
KM	0.79	0.77
tSNE-KM	0.63	0.73
PCA-KM	0.81	0.79
NMF-KM	0.77	0.78
SKM	0.15	0.23
MClust	0.83	0.79

Table A2: ARIs and NMIs of different methods on scRNA-Seq data. The number of clusters is set to be truth ($K = 8$) for all methods.

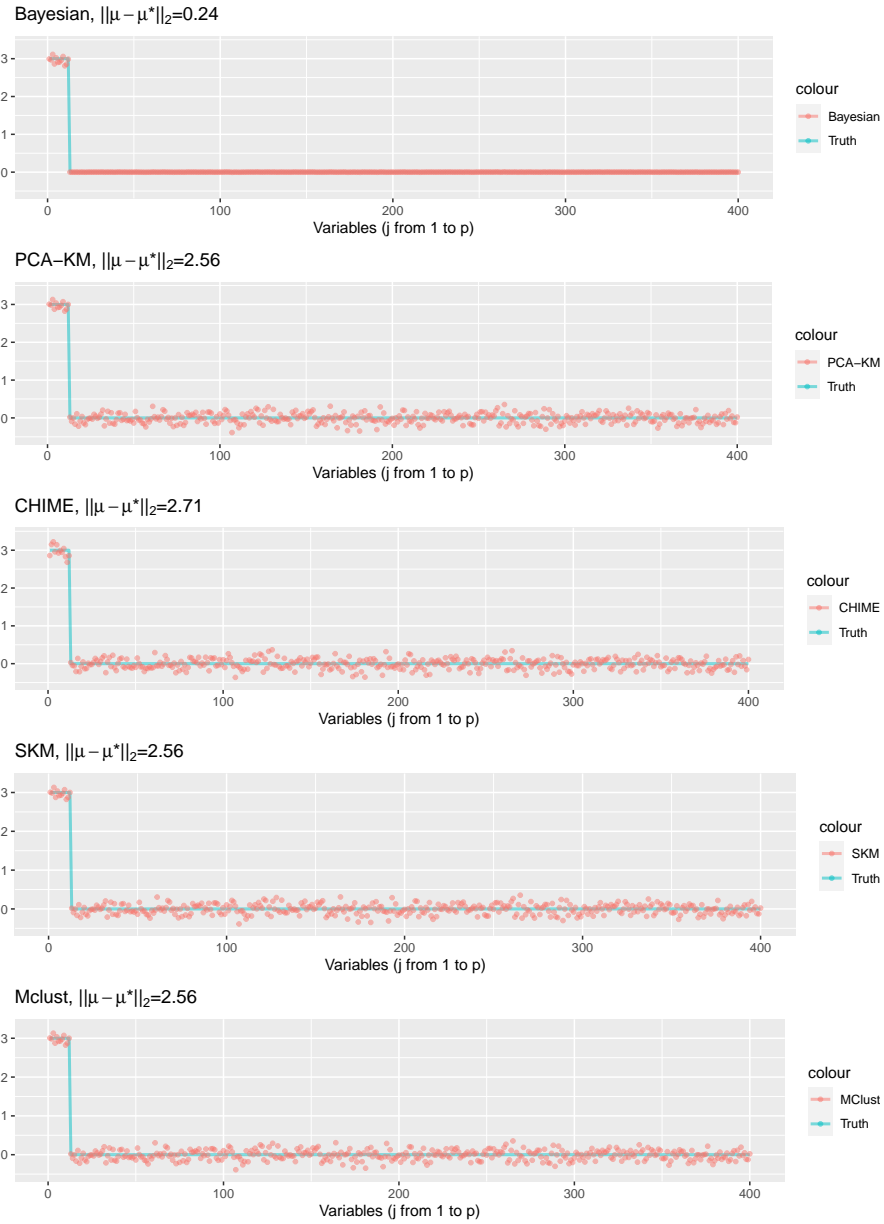


Figure A2: Estimation of mean vector μ_1^* of different methods in Scenario I with $K^* = 3$ and $s = 12$ in a randomly selected simulation replicate.

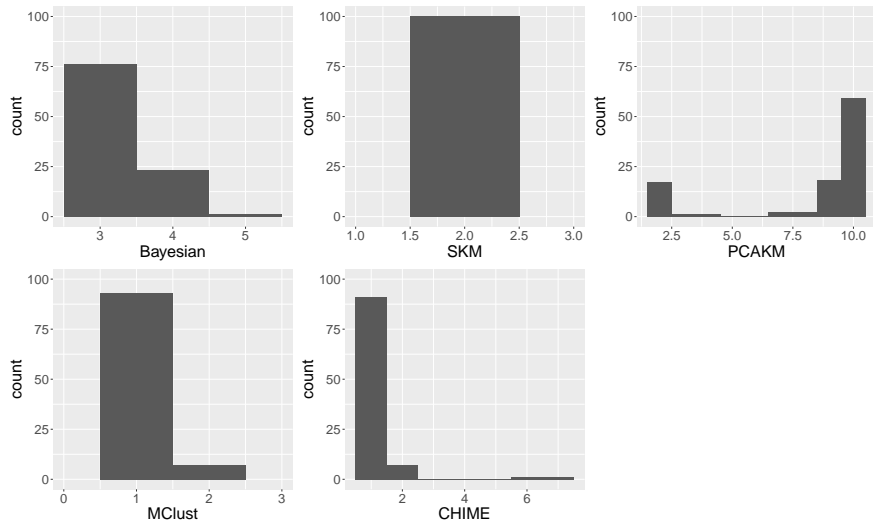


Figure A3: Histograms of estimated number of clusters under different methods in Scenario III of simulation studies.

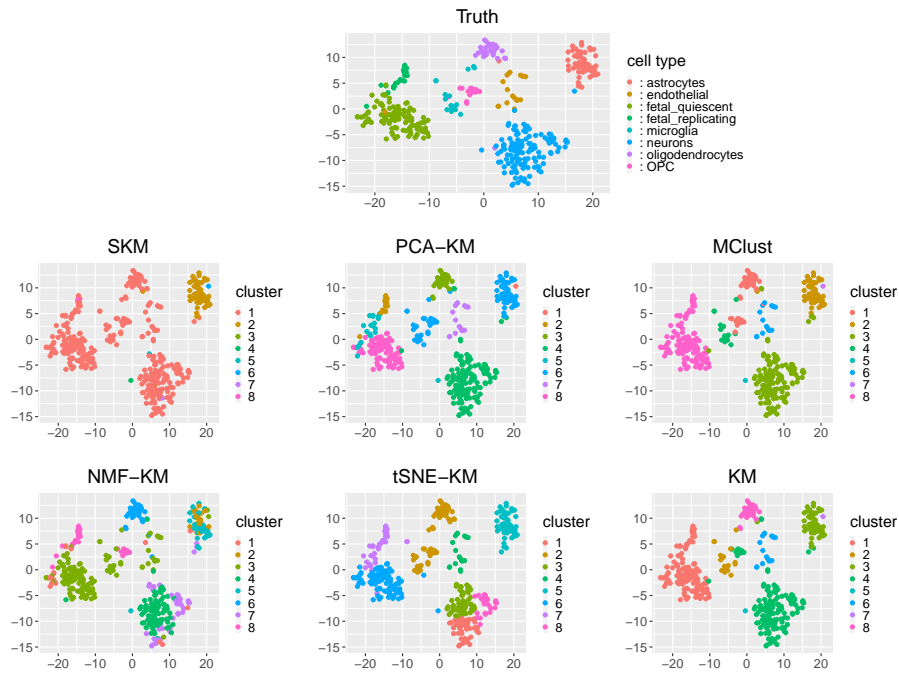


Figure A4: Clustering results of scRNA-Seq data corresponding to alternative methods. The number of clusters is set to be truth ($K = 8$) for all methods.

References

- Joshua Agterberg, Zachary Lubberts, and Carey E Priebe. Entrywise estimation of singular vectors of low-rank matrices with heteroskedasticity and dependence. *IEEE Transactions on Information Theory*, 68(7):4618–4650, 2022.
- Hassan Ashtiani, Shai Ben-David, Nicholas JA Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Near-optimal sample complexity bounds for robust learning of gaussian mixtures via compression schemes. *Journal of the ACM (JACM)*, 67(6):1–42, 2020.
- Martin Azizyan, Aarti Singh, and Larry Wasserman. Minimax theory for high-dimensional gaussian mixtures with sparse mean separation. *Advances in Neural Information Processing Systems*, 26, 2013.
- Ray Bai, Veronika Ročková, and Edward I George. Spike-and-slab meets lasso: A review of the spike-and-slab lasso. *Handbook of Bayesian Variable Selection*, pages 81–108, 2021.
- Anirban Bhattacharya, Debdeep Pati, Natesh S. Pillai, and David B. Dunson. Dirichlet–Laplace priors for optimal shrinkage. *Journal of the American Statistical Association*, 110(512):1479–1490, 2015. doi: 10.1080/01621459.2014.960967. URL <https://doi.org/10.1080/01621459.2014.960967>. PMID: 27019543.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- T. Tony Cai, Jing Ma, and Linjun Zhang. CHIME: Clustering of high-dimensional Gaussian mixtures with EM algorithm and its optimality. *The Annals of Statistics*, 47(3):1234 – 1267, 2019. doi: 10.1214/18-AOS1711. URL <https://doi.org/10.1214/18-AOS1711>.
- Junyue Cao, Malte Spielmann, Xiaojie Qiu, Xingfan Huang, Daniel M Ibrahim, Andrew J Hill, Fan Zhang, Stefan Mundlos, Lena Christiansen, Frank J Steemers, et al. The single-cell transcriptional landscape of mammalian organogenesis. *Nature*, 566(7745):496–502, 2019.
- Joshua Cape, Minh Tang, and Carey Priebe. Signal-plus-noise matrix models: Eigenvector deviations and fluctuations. *Biometrika*, 106, 02 2018. doi: 10.1093/biomet/asy070.
- Carlos Carvalho, Nicholas Polson, and James Scott. Handling sparsity via the Horseshoe. *Journal of Machine Learning Research - Proceedings Track*, 5:73–80, 01 2009.
- Ismaël Castillo and Aad van der Vaart. Needles and Straw in a Haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics*, 40(4):2069 – 2101, 2012. doi: 10.1214/12-AOS1029. URL <https://doi.org/10.1214/12-AOS1029>.
- Ismaël Castillo, Johannes Schmidt-Hieber, and Aad van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986 – 2018, 2015. doi: 10.1214/15-AOS1334. URL <https://doi.org/10.1214/15-AOS1334>.

- Noirrit Kiran Chandra, Antonio Canale, and David B Dunson. Escaping the curse of dimensionality in Bayesian model based clustering. *arXiv preprint arXiv:2006.02700*, 2022.
- Spyros Darmanis, Steven A Sloan, Ye Zhang, Martin Enge, Christine Caneda, Lawrence M Shuer, Melanie G Hayden Gephart, Ben A Barres, and Stephen R Quake. A survey of human brain transcriptome diversity at the single cell level. *Proceedings of the National Academy of Sciences*, 112(23):7285–7290, 2015.
- E Diday, G Govaert, Y Lechevallier, and J Sidi. Clustering in pattern recognition. In *Digital Image Processing*, pages 19–58. Springer, 1981.
- Chris Ding, Xiaofeng He, Hongyuan Zha, and Horst D Simon. Adaptive dimension reduction for clustering high dimensional data. In *2002 IEEE International Conference on Data Mining, 2002. Proceedings.*, pages 147–154. IEEE, 2002.
- Natalie Doss, Yihong Wu, Pengkun Yang, and Harrison H Zhou. Optimal estimation of high-dimensional gaussian location mixtures. *The Annals of Statistics*, 51(1):62–95, 2023.
- Chris Fraley, Adrian Raftery, Thomas Murphy, and Luca Scrucca. MCLUST Version 4 for R: Normal Mixture Modeling for Model-Based Clustering, Classification, and Density Estimation. *Technical Report No. 597*, 01 2012.
- Chao Gao, Aad W van der Vaart, and Harrison H Zhou. A general framework for bayes structured linear models. 2020.
- Subhashis Ghosal, Jayanta K. Ghosh, and Aad W. van der Vaart. Convergence rates of posterior distributions. *The Annals of Statistics*, 28(2):500 – 531, 2000. doi: 10.1214/aos/1016218228. URL <https://doi.org/10.1214/aos/1016218228>.
- Joydeep Ghosh and Ayan Acharya. Cluster ensembles. *Wiley interdisciplinary reviews: Data mining and knowledge discovery*, 1(4):305–315, 2011.
- Prasenjit Ghosh, Debdeep Pati, and Anirban Bhattacharya. Posterior contraction rates for stochastic block models. *Sankhya A*, 82(2):448–476, 2020.
- Jiajun Gu and Jun S Liu. Bayesian biclustering of gene expression data. *BMC genomics*, 9(1):1–10, 2008.
- Torben Hagerup and Christine Rüb. A guided tour of chernoff bounds. *Information processing letters*, 33(6):305–308, 1990.
- Jiashun Jin and Wanjie Wang. Influential features PCA for high dimensional clustering. *The Annals of Statistics*, 44(6):2323 – 2359, 2016.
- Jiashun Jin, Zheng Tracy Ke, and Wanjie Wang. Phase transitions for high dimensional clustering and related problems. *The Annals of Statistics*, 45(5):2151–2189, 2017. ISSN 00905364. URL <http://www.jstor.org/stable/26362898>.
- Michael Kosorok. *Introduction to Empirical Processes and Semiparametric Inference*. 01 2008. ISBN 978-0-387-74977-8. doi: 10.1007/978-0-387-74978-5.

- Tianyang Li, Xinyang Yi, Constantine Carmanis, and Pradeep Ravikumar. Minimax Gaussian classification & clustering. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR, 2017.
- George C Linderman, Manas Rachh, Jeremy G Hoskins, Stefan Steinerberger, and Yuval Kluger. Fast interpolation-based t-SNE for improved visualization of single-cell RNA-seq data. *Nature methods*, 16(3):243–245, 2019.
- Matthias Löffler, Anderson Y Zhang, and Harrison H Zhou. Optimality of spectral clustering in the Gaussian mixture model. *The Annals of Statistics*, 49(5):2506–2530, 2021.
- Ulrike Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17:395–416, 01 2004. doi: 10.1007/s11222-007-9033-z.
- Pascal Massart. Concentration inequalities and model selection. école d’été de probabilités de saint-flour xxxiii – 2003. *Lecture Notes in Mathematics -Springer-verlag-*, 1896, 01 2007. doi: 10.1007/978-3-540-48503-2.
- Jeffrey W Miller. Consistency of mixture models with a prior on the number of components. *Dependence Modeling*, 11(1):20220150, 2023.
- Jeffrey W. Miller and Matthew T. Harrison. Mixture models with a prior on the number of components. *Journal of the American Statistical Association*, 113(521):340–356, 2018. doi: 10.1080/01621459.2016.1255636. URL <https://doi.org/10.1080/01621459.2016.1255636>.
- T. J. Mitchell and J. J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988. ISSN 01621459. URL <http://www.jstor.org/stable/2290129>.
- Pierpaolo Natalini and Biagio Palumbo. Inequalities for the incomplete gamma function. *Math. Inequal. Appl*, 3(1):69–77, 2000.
- Ilsang Ohn and Lizhen Lin. Optimal Bayesian estimation of Gaussian mixtures with growing number of components. *arXiv preprint arXiv:2007.09284*, 2022.
- Debdeep Pati, Anirban Bhattacharya, Natesh S Pillai, and David Dunson. Posterior contraction in sparse bayesian factor models for massive covariance matrices. 2014.
- William M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850, 1971. ISSN 01621459. URL <http://www.jstor.org/stable/2284239>.
- Kolyan Ray and Botond Szabó. Variational bayes for high-dimensional linear regression with sparse priors. *Journal of the American Statistical Association*, pages 1–12, 2021.
- Veronika Ročková and Edward I George. The spike-and-slab lasso. *Journal of the American Statistical Association*, 113(521):431–444, 2018.
- Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.

- Padhraic Smyth. Model selection for probabilistic clustering using cross-validated likelihood. *Statistics and computing*, 10(1):63–72, 2000. doi: 10.1023/A:1008940618127.
- Mahlet G Tadesse and Marina Vannucci. Handbook of Bayesian variable selection. 2021.
- Mahlet G Tadesse, Naijun Sha, and Marina Vannucci. Bayesian variable selection in clustering high-dimensional data. *Journal of the American Statistical Association*, 100(470):602–617, 2005.
- Robert Tibshirani, Guenther Walther, and Trevor Hastie. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):411–423, 2001.
- Daniela M. Witten and Robert Tibshirani. A framework for feature selection in clustering. *Journal of the American Statistical Association*, 105(490):713–726, 2010. ISSN 01621459. URL <http://www.jstor.org/stable/29747077>.
- Fangzheng Xie, Joshua Cape, Carey E Priebe, and Yanxun Xu. Bayesian sparse spiked covariance model with a continuous matrix shrinkage prior. *Bayesian Analysis*, 1(1):1–25, 2022.
- Dan Yang, Zongming Ma, and Andreas Buja. Rate optimal denoising of simultaneously sparse and low rank matrices. *Journal of Machine Learning Research*, 17(92):1–27, 2016. URL <http://jmlr.org/papers/v17/15-134.html>.
- Bin Yu. *Assouad, Fano, and Le Cam*, pages 423–435. Springer New York, New York, NY, 1997. ISBN 978-1-4612-1880-7. doi: 10.1007/978-1-4612-1880-7_29. URL https://doi.org/10.1007/978-1-4612-1880-7_29.
- Y. Yu, T. Wang, and R. J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 04 2014. ISSN 0006-3444. doi: 10.1093/biomet/asv008. URL <https://doi.org/10.1093/biomet/asv008>.
- Xun Zhu, Travers Ching, Xinghua Pan, Sherman M Weissman, and Lana Garmire. Detecting heterogeneity in single-cell RNA-Seq data by non-negative matrix factorization. *PeerJ*, 5:e2888, 2017.