ERROR BOUNDS FOR SPECTRAL CLUSTERING OVER SAMPLES FROM SPHERICAL GAUSSIAN MIXTURE MODELS

Zhaoqiang Liu

Department of ECE, National University of Singapore

ABSTRACT

Spectral clustering has been one of the most popular methods for clustering multivariate data and has been widely used in image processing and data mining. Despite its considerable empirical success, the theoretical properties of spectral clustering are not yet fully developed. In this paper, we derive upper bounds for the clustering error of spectral clustering for data samples generated from spherical Gaussian mixture models. In our analysis, first, the graph Laplacian calculated from samples is approximated by a reference graph Laplacian which has good spectral properties. Second, we use the Davis-Kahan perturbation theorem to provide an upper bound for the sum of squared distances between each projected data point and its cluster center. Finally, we leverage theoretical results of Meilă's to prove an upper bound for the clustering error from the upper bound for the sum of squared distances.

Index Terms— Spectral clustering, Gaussian Mixture Models, Optimal clusterings, Error bounds

1. INTRODUCTION

Clustering is a fundamental and ubiquitous problem in various applications. Spectral clustering has become one of the most widely-used clustering algorithms. It has been shown to outperform various classic clustering algorithms such as k-means [1] on a number of challenging clustering problems. Suppose we have a data matrix of N samples $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N] \in \mathbb{R}^{F\times N}$ and a pre-specified number of clusters $K < N$, the spectral clustering algorithm first constructs a projected data matrix $\hat{\mathbf{V}} = [\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_N] \in \mathbb{R}^{K\times N}$, where the rows of $\hat{\mathbf{V}}$ correspond to eigenvectors of the K smallest eigenvalues of the graph Laplacian matrix. Then the projected data points are clustered into k clusters via k-means (or other proper clustering algorithms). For a data matrix with N samples, a K*-clustering* (or simply a *clustering*) is defined as a set of pairwise disjoint index sets $\mathscr{C} := \{\mathscr{C}_1, \mathscr{C}_2, \dots, \mathscr{C}_K\}$ whose union is $\{1, 2, \dots, N\}$. The *sum-of-squares distortion measure* with respect to the data matrix $\hat{\mathbf{V}}$ and a K-clustering $\mathscr{C}$ is defined as

$$\mathcal{D}(\hat{\mathbf{V}}, \mathscr{C}) := \sum_{k=1}^{K} \sum_{n\in\mathscr{C}_k} \|\hat{\mathbf{v}}_n - \hat{\mathbf{c}}_k\|_2^2, \tag{1}$$

where $\hat{\mathbf{c}}_k := \frac{1}{|\mathscr{C}_k|}\sum_{n\in\mathscr{C}_k} \hat{\mathbf{v}}_n$ is the cluster center. The k-means algorithm over $\hat{\mathbf{V}}$ tries to find an *optimal clustering* $\mathscr{C}^{\mathrm{opt}}$ that satisfies

$$\mathcal{D}(\hat{\mathbf{V}}, \mathscr{C}^{\mathrm{opt}}) = \min_{\mathscr{C}} \mathcal{D}(\hat{\mathbf{V}}, \mathscr{C}), \tag{2}$$

where the minimization is taken over all K-clusterings. For brevity, we say such $\mathscr{C}^{\mathrm{opt}}$ is an optimal clustering of spectral clustering (over $\mathbf{V}$). Although spectral clustering has enjoyed success in wide-ranging practical areas, there has been little work on the theoretical analysis of it [2–9].

A K-component Gaussian mixture model (GMM) is a generative model that assumes there are K multivariate Gaussian distributions and a probability vector $\mathbf{w} := [w_1, w_2, \dots, w_K]$, such that data samples are independently sampled and the probability that each sample comes from the k-th component is w_k. w_k is said to be the *mixing weight for the k-th component*, and we use $\mathbf{u}_k, \mathbf{\Sigma}_k$ to denote the *component mean vector* and the *component covariance matrix* respectively. When $\mathbf{\Sigma}_k = \sigma_k^2\mathbf{I}$, where $\mathbf{I}$ is the identity matrix, we say the GMM is *spherical*. Suppose that $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N$ are samples independently generated from a K-component GMM, the *correct target clustering* $\mathscr{C} := \{\mathscr{C}_1, \mathscr{C}_2, \dots, \mathscr{C}_K\}$ satisfies the condition that $n \in \mathscr{C}_k$ if and only if $\mathbf{v}_n$ comes from the k-th component.

1.1. Main Contributions

We prove if the samples are independently generated from a K-component spherical GMM with an appropriate separability assumption, then the clustering error of spectral clustering (with an unnormalized graph Laplacian) can be upper bounded (by bounding the distance between any optimal clustering and the correct target clustering) with high probability when the number of samples is sufficiently large. We mention in Section 5 that such theoretical results can be extended to spectral clustering with normalized graph Laplacians.

1.2. Notations

We use a_{ij} or $[\mathbf{A}]_{ij}$ to denote the (i,j)-th entry of $\mathbf{A}$. $[N]$ represents $\{1, 2, \cdots, N\}$ for any positive integer N.

The work of Z. Liu (a0109983@u.nus.edu, zqliu12@gmail.com) is supported by NUS IDS Grant R-263-000-C73-646.

$\|\mathbf{V}\|_2, \|\mathbf{V}\|_\mathrm{F}$ represent the spectral norm and the Frobenius norm of $\mathbf{V}$ respectively. $\mathrm{Diag}(\mathbf{w})$ represents the diagonal matrix whose diagonal entries are given by $\mathbf{w}$. The eigendecomposition of a *positive semi-definite* (PSD) matrix $\mathbf{L} \in \mathbb{R}^{N\times N}$ is given by $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T$ with $\mathbf{U} \in \mathbb{R}^{N\times N}$ being an orthogonal matrix and $\mathbf{\Sigma} \in \mathbb{R}^{N\times N}$ being a diagonal matrix of eigenvalues. $\lambda_n(\mathbf{L})$ represents the n-th largest eigenvalue of $\mathbf{L}$ and the diagonal entries of $\mathbf{\Sigma}$ are ordered such that $[\mathbf{\Sigma}]_{nn} = \lambda_n(\mathbf{L})$. For any $n \in [N]$, we say that $\mathbf{u}_n := \mathbf{U}(:, n)$ is the n-th eigenvector of $\mathbf{L}$.

2. RELATED WORK

2.1. Spectral Clustering

Given a data matrix $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_N] \in \mathbb{R}^{F\times N}$, we obtain a similarity matrix $\mathbf{A} \in \mathbb{R}^{N\times N}$ by setting

$$a_{ij} = \exp\left(-\phi\|\mathbf{v}_i - \mathbf{v}_j\|_2^2\right) \text{ for } i, j \in [N], \tag{3}$$

where $\phi > 0$ is a scaling parameter. Define $\mathbf{D}$ to be the diagonal matrix in $\mathbb{R}^{N\times N}$ such that $d_{nn} = \sum_{m=1}^N a_{nm}$ for $n \in [N]$. The (unnormalized) graph Laplacian $\mathbf{L}$ is constructed as $\mathbf{L} := \mathbf{D} - \mathbf{A}$. According to [10], $\mathbf{L}$ is PSD. Let the eigendecomposition (cf. Section 1.2) of $\mathbf{L}$ be $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T$. We have that $\lambda_N(\mathbf{L}) = 0$ and $\mathbf{u}_N = \frac{\mathbf{e}}{\sqrt{N}}$, where $\mathbf{e}$ is the vector of all ones. The spectral clustering with K clusters first derives the projected data matrix $\hat{\mathbf{V}} := \mathbf{U}(:, (N-K+1)\colon N)^T \in \mathbb{R}^{K\times N}$. Then the k-means algorithm is performed on $\hat{\mathbf{V}}$ to obtain an (approximately) optimal clustering.

2.2. Lower Bound on Distortion and the ME Distance

Suppose we have a dataset $\hat{\mathbf{V}} = [\hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N] \in \mathbb{R}^{K\times N}$ and a K-clustering $\mathscr{C} := \{\mathscr{C}_1, \mathscr{C}_2, \ldots, \mathscr{C}_K\}$. Let $\mathbf{H} \in \mathbb{R}^{N\times K}$ be the binary matrix satisfying $h_{nk} = 1$ if and only if $n \in \mathscr{C}_k$. Let $n_k = |\mathscr{C}_k|$, and $\bar{\mathbf{H}} := \mathbf{H}\mathrm{Diag}(\frac{1}{\sqrt{n_1}}, \frac{1}{\sqrt{n_2}}, \ldots, \frac{1}{\sqrt{n_K}})$. We have $\bar{\mathbf{H}}^T\bar{\mathbf{H}} = \mathbf{I}$ and the distortion can be written as [11]

$$\mathcal{D}(\hat{\mathbf{V}}, \mathscr{C}) = \|\hat{\mathbf{V}} - \hat{\mathbf{V}}\bar{\mathbf{H}}\bar{\mathbf{H}}^T\|_\mathrm{F}^2. \tag{4}$$

Let $\hat{\mathbf{Z}}$ be the centralized data matrix of $\hat{\mathbf{V}}$ and define $\hat{\mathbf{S}} := \hat{\mathbf{Z}}^T\hat{\mathbf{Z}}$. For any two K-clusterings, we define the so-called *misclassification error (ME) distance* to compare their structures.

Definition 1 *Let $\mathcal{P}_K$ be the set of all permutations of $[K]$. The misclassification error distance of any two K-clusterings $\mathscr{C}^1 := \{\mathscr{C}_1^1, \mathscr{C}_2^1, \ldots, \mathscr{C}_K^1\}$ and $\mathscr{C}^2 := \{\mathscr{C}_1^2, \mathscr{C}_2^2, \ldots, \mathscr{C}_K^2\}$ is $\mathrm{d_{ME}}(\mathscr{C}^1, \mathscr{C}^2) := 1 - \frac{1}{N}\max_{\pi\in\mathcal{P}_K}\sum_{k=1}^K |\mathscr{C}_k^1 \bigcap \mathscr{C}_{\pi(k)}^2|$.*

From [12], we know that the ME distance defined above is indeed a distance. Define $\tau(\delta) := 2\delta(1 - \delta/(K-1))$. We have the following lemma [13] which enables us to obtain an upper bound for the ME distance between a clustering and any optimal clustering provided the distortion of this clustering is appropriately upper bounded.

Lemma 1 *$\mathscr{C} := \{\mathscr{C}_1, \ldots, \mathscr{C}_K\}$ is a K-clustering of $\hat{\mathbf{V}} \in \mathbb{R}^{K\times N}$ and $p_{\max} := \max_k \frac{|\mathscr{C}_k|}{N}$, $p_{\min} := \min_k \frac{|\mathscr{C}_k|}{N}$. Define $\delta := \frac{\mathcal{D}(\hat{\mathbf{V}}, \mathscr{C}) - \lambda_K(\hat{\mathbf{S}})}{\lambda_{K-1}(\hat{\mathbf{S}}) - \lambda_K(\hat{\mathbf{S}})}$. Then if $\delta \le \frac{1}{2}(K-1)$ and $\tau(\delta) \le p_{\min}$,*

$$\mathrm{d_{ME}}(\mathscr{C}, \mathscr{C}^{\mathrm{opt}}) \le p_{\max}\tau(\delta), \tag{5}$$

where $\mathscr{C}^{\mathrm{opt}}$ is an optimal K-clustering for $\hat{\mathbf{V}}$.

3. THE MAIN THEOREM

Let $Y = e^{-X^2}$, where X follows the Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$, we have $\mathbb{E}[Y] = \frac{1}{\sqrt{2\sigma^2+1}}\exp\left(-\frac{\mu^2}{2\sigma^2+1}\right)$, and $\mathrm{Var}[Y] = \frac{1}{\sqrt{4\sigma^2+1}}\exp\left(-\frac{2\mu^2}{4\sigma^2+1}\right) - (\mathbb{E}[Y])^2$. Let

$$s_{ij} := \left(2\phi(\sigma_i^2 + \sigma_j^2) + 1\right)^{-\frac{F}{2}}\exp\left(-\frac{\phi\|\mathbf{u}_i - \mathbf{u}_j\|_2^2}{2\phi(\sigma_i^2 + \sigma_j^2) + 1}\right),$$

$$t_{ij} := \left(4\phi(\sigma_i^2 + \sigma_j^2) + 1\right)^{-\frac{F}{2}}\exp\left(-\frac{2\phi\|\mathbf{u}_i - \mathbf{u}_j\|_2^2}{4\phi(\sigma_i^2 + \sigma_j^2) + 1}\right) - s_{ij}^2$$

for $i, j \in [K]$. Let $T = \max_{i,j} t_{ij}$ and $S_o = \max_{i\ne j} s_{ij}$. For $p \in [0, \frac{K-1}{2}]$, we define the function $\zeta(p) := \frac{p}{1+\sqrt{1-2p/(K-1)}}$. Let $w_{\min} := \min_k w_k$, $w_{\max} := \max_k w_k$ and $c_k = w_k s_{kk} + \sum_{k'\ne k} w_{k'} s_{k,k'}$. Our main theorem is as follows:

Theorem 2 *Suppose the columns of $\mathbf{V} \in \mathbb{R}^{F\times N}$ are independently generated from a K-component spherical GMM and $N > K$. Assume that there is a $\xi \in (0,1)$, such that $\min_k c_k - \sqrt{2K}S_o > 2\sqrt{T}/\sqrt{\xi}$. Further assume*

$$\frac{4(K-1)T/\xi}{\left(\min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi}\right)^2} < \zeta(w_{\min}). \tag{6}$$

Let $\mathscr{C} := \{\mathscr{C}_1, \mathscr{C}_2, \ldots, \mathscr{C}_K\}$ be the correct target K-clustering corresponding to the spherical GMM. Assume that $\epsilon > 0$ is sufficiently small.[1] *Then with probability at least $1 - \xi - Ke\exp\left(-cN\epsilon^2\right)$, $\mathrm{d_{ME}}(\mathscr{C}, \mathscr{C}^{\mathrm{opt}})$ is upper bounded by*

$$\tau\left(\frac{4(K-1)T/\xi}{\left(\min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi} - \epsilon\right)^2}\right)(w_{\max} + \epsilon), \tag{7}$$

where $c > 0$ is a constant depending on $\{w_k\}_{k\in[K]}$, and $\mathscr{C}^{\mathrm{opt}}$ is an optimal K-clustering of spectral clustering over $\mathbf{V}$.

Remark 1 *Considering the special case that $\sigma_1^2 = \ldots = \sigma_K^2 = \sigma^2$ and $\|\mathbf{u}_i - \mathbf{u}_j\|_2 = d$ for all $1 \le i \ne j \le K$. Choosing $\phi > 0$ to be small such that $\phi d^2 \ll 1$ and $F\phi\sigma^2 \ll$*

[1] $\epsilon \le \min\left\{\frac{w_{\min}}{2}, \min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi}\right\}$ and ϵ is chosen such that $\frac{4(K-1)T/\xi}{\left(\min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi} - \epsilon\right)^2} < \zeta(w_{\min} - \epsilon)$.

1. *We show (cf. Section 4.5 to follow) that for this case, the separability assumption* (6) *can be expressed as*

$$d > \frac{2\sigma}{\sqrt{w_{\min}}}\left(\frac{2F}{\xi}\right)^{1/4}\left[1+\left(\frac{2(K-1)}{w_{\min}}\right)^{1/4}\right]. \quad (8)$$

Remark 2 *The k-means algorithm often gets stuck at local optima and is not able to find a globally optimal solution. For this issue, note that our theoretical results can be easily extended to provide upper bounds for the ME distance between any* ***approximately optimal*** *clustering and the correct target clustering. See Section 5.1 in [14] for a detailed discussion.*

4. PROOF OF THE MAIN THEOREM

Before proving the theorem, we prove the following lemmas.

4.1. Bounding the Distance between the Empirical and Reference Graph Laplacians

First, we construct a reference graph Laplacian matrix $\bar{\mathbf{L}}$ and derive an upper bound for the spectral norm of $\mathbf{L}-\bar{\mathbf{L}}$. Let the reference similarity matrix $\bar{\mathbf{A}} \in \mathbb{R}^{N\times N}$ satisfy that

$$\bar{a}_{n_1,n_2} = s_{ij} \text{ if } n_1 \in \mathscr{C}_i, n_2 \in \mathscr{C}_j, \quad (9)$$

for $n_1, n_2 \in [N]$, $n_1 \neq n_2$. We set $\bar{a}_{nn} = 1$ for all $n \in [N]$. Constructing the diagonal matrix $\bar{\mathbf{D}} \in \mathbb{R}^{N\times N}$ such that $\bar{d}_{nn} = \sum_{m=1}^{N} \bar{a}_{nm}$ and set $\bar{\mathbf{L}} = \bar{\mathbf{D}} - \bar{\mathbf{A}}$. Let $n_k = |\mathscr{C}_k|$ for $k \in [K]$ and set $\mathbf{E} = \mathbf{L} - \bar{\mathbf{L}}$, we have the following lemma which bounds the distance between $\mathbf{L}$ and $\bar{\mathbf{L}}$.

Lemma 3 *For any $\xi \in (0,1)$, with probability at least $1-\xi$,*

$$\|\mathbf{E}\|_2 < \frac{2N\sqrt{T}}{\sqrt{\xi}}. \quad (10)$$

Proof Splitting $\mathbf{E}$ as $\mathbf{E} = \mathbf{E}_D + \mathbf{E}_O$, where $\mathbf{E}_D$ and $\mathbf{E}_O$ are the diagonal and off-diagonal parts of $\mathbf{E}$ respectively. Note that $\mathbb{E}[e_{ij}] = 0$ for $i,j \in [N]$. We have

$$\begin{aligned}
&\mathbb{E}\left[\|\mathbf{E}\|_2^2\right] \le 2\left[\mathbb{E}\left[\|\mathbf{E}_D\|_2^2\right] + \mathbb{E}\left[\|\mathbf{E}_O\|_2^2\right]\right] && (11)\\
&\le 2\mathbb{E}\left[\|\mathbf{E}_D\|_2^2\right] + 2\mathbb{E}\left[\|\mathbf{E}_O\|_{\mathrm{F}}^2\right] && (12)\\
&= 2\max_{n\in[N]} \mathbb{E}[e_{nn}^2] + 2\sum_{i\neq j}\mathbb{E}\left[e_{ij}^2\right] && (13)\\
&= 2\max_{n\in[N]} \mathrm{Var}[e_{nn}] + 2\sum_{i\neq j}\mathrm{Var}[e_{ij}] && (14)\\
&= 2\max_{n\in[N]} \mathrm{Var}[e_{nn}] + 2\sum_{i\neq j}\mathrm{Var}[a_{ij}] && (15)\\
&\le 2(N-1)^2T + 2\sum_{k=1}^{K} n_k\left((n_k-1)t_{kk} + \sum_{k'\neq k} n_{k'}t_{k,k'}\right)\\
&\le 4N^2T. && (16)
\end{aligned}$$

By Markov inequality, we have that for any $\xi \in (0,1)$,

$$\mathbb{P}(\|\mathbf{E}\|_2^2 \ge \frac{1}{\xi}\mathbb{E}\left[\|\mathbf{E}\|_2^2\right]) \le \xi, \quad (17)$$

which completes the proof. ∎

Remark 3 *Note that it is important to split $\mathbf{E}$ into diagonal and off-diagonal parts. Otherwise, we will derive an upper bound for $\mathbb{E}\left[\|\mathbf{E}\|_2\right]$ from the upper bound of $\mathbb{E}\left[\|\mathbf{E}\|_{\mathrm{F}}\right]$, which is $O(N^{3/2})$, and is disastrously loose for further derivations.*

4.2. Spectral Properties of $\bar{\mathbf{L}}$

Next, we describe the spectral properties of $\bar{\mathbf{L}}$.

Lemma 4 *For any $k \in [K]$, let $g_k := \frac{n_k s_k + \sum_{k'\neq k} n_{k'} s_{k,k'}}{N}$. We have that Ng_k is an eigenvalue of $\bar{\mathbf{L}}$ with multiplicity (at least) $n_k - 1$. Moreover, if $\sqrt{2K}S_o < \min_k g_k$, we have that $\bar{\lambda}_{N-K+1} \le \sqrt{2K}NS_o$ and for $N-K < n < N$, the eigenvector $\bar{\mathbf{u}}_n$ satisfies that $\bar{u}_n(i) = \bar{u}_n(j)$ if the i-th and the j-th samples are generated from the same Gaussian component.*

Proof We have that if $g_1, \ldots, g_K$ are all distinct, then

$$\begin{aligned}
&\{\mathbf{x} : (\bar{\mathbf{L}} - Ng_k\mathbf{I})\mathbf{x} = \mathbf{0}\}\\
&= \{\mathbf{x} : \sum_{n\in\mathscr{C}_k} x_n = 0, x_n = 0 \text{ for } n \notin \mathscr{C}_k\}. \quad (18)
\end{aligned}$$

Therefore, Ng_k is an eigenvalue of $\bar{\mathbf{L}}$ with multiplicity $n_k - 1$. If the g_k's are not all distinct. For any $g \in \{g_k\}_{k\in[K]}$. Let $k_g := \{k \in [K] : g_k = g\}$. Then the multiplicity of eigenvalue Ng is $\sum_{k\in k_g}(n_k - 1)$. On the other hand, if we consider the eigenvector $\mathbf{x}$ in the form that $x_n = \alpha_k$ if $n \in \mathscr{C}_k$. Then the linear equation system $\bar{\mathbf{L}}\mathbf{x} = \lambda\mathbf{x}$ can be reformulated as

$$\mathbf{F}\boldsymbol{\alpha} = \lambda\boldsymbol{\alpha}, \quad (19)$$

where λ is the corresponding eigenvalue, $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \ldots, \alpha_K]^T$ and $\mathbf{F} \in \mathbb{R}^{K\times K}$ satisfies that

$$f_{k,k'} = -n_{k'}s_{k,k'} \text{ if } k \neq k', \quad f_{kk} = \sum_{k'\neq k} n_{k'}s_{k,k'}. \quad (20)$$

Similar to that for $\bar{\mathbf{L}}$, we have that $\mathbf{F}$ is PSD with the smallest eigenvalue being 0. In addition, all the eigenvalues of $\mathbf{F}$ are also eigenvalues of $\bar{\mathbf{L}}$. Splitting $\mathbf{F}$ as $\mathbf{F} = \mathbf{F}_D + \mathbf{F}_O$, similar to that in the proof of Lemma 3, we derive that $\|\mathbf{F}\|_2^2 \le 2KN^2S_o^2$. Therefore, if $\sqrt{2K}S_o < \min_k g_k$, we have that $\lambda_{N-K+1}(\bar{\mathbf{L}}) = \|\mathbf{F}\|_2 \le \sqrt{2K}NS_o$ and for $N-K < n < N$, $\bar{\mathbf{u}}_n$ satisfies that $\bar{u}_n(i) = \bar{u}_n(j)$ if the i-th and the j-th samples come from the same Gaussian component. ∎

Remark 4 *A special case is that for any $k \neq k'$, $s_{k,k'} = s$ (e.g., when $\sigma_1^2 = \ldots = \sigma_K^2 = \sigma^2$ and $\|\mathbf{u}_i - \mathbf{u}_j\|_2 = d$ for all $i \neq j$). For this case, $\lambda_{N-1}(\bar{\mathbf{L}}) = \ldots = \lambda_{N-K+1}(\bar{\mathbf{L}}) = Ns$.*

4.3. Bounding the Distortion for Projected Data Matrix

Let the eigendecomposition of $\bar{\mathbf{L}}$ be $\bar{\mathbf{L}} = \bar{\mathbf{U}}\bar{\mathbf{\Sigma}}\bar{\mathbf{U}}^T$. In addition, let $\bar{\mathbf{B}} = \bar{\mathbf{U}}(:, (N-K+1):N)$, $\mathbf{B} = \mathbf{U}(:, (N-K+1):N)$ (i.e., $\mathbf{B}^T = \hat{\mathbf{V}}$), $\bar{\mathbf{\Lambda}} = \bar{\mathbf{\Sigma}}((N-K+1):N, (N-K+1):N)$ and $\mathbf{\Lambda} = \mathbf{\Sigma}((N-K+1):N, (N-K+1):N)$. Let

$\bar{\mathbf{U}}_{N-K} = \bar{\mathbf{U}}(:,1:(N-K))$, $\bar{\mathbf{\Sigma}}_{N-K} = \bar{\mathbf{\Sigma}}(1:(N-K),1:(N-K))$, $\mathbf{U}_{N-K} = \mathbf{U}(:,1:(N-K))$ and $\mathbf{\Sigma}_{N-K} = \mathbf{\Sigma}(1:(N-K),1:(N-K))$. In the following lemma, we provide an upper bound for the distortion $\mathcal{D}(\hat{\mathbf{V}},\mathscr{C})$ using the Davis-Kahan perturbation theorem [15].

Lemma 5 *If* $\eta := N\min_k g_k - (\sqrt{2K}NS_o + \|\mathbf{E}\|_2) > 0$, *then we have that*

$$\mathcal{D}(\hat{\mathbf{V}},\mathscr{C}) \leq \frac{(K-1)\|\mathbf{E}\|_2^2}{\eta^2}. \tag{21}$$

Proof By (19), the subspaces spanned by the columns of $\bar{\mathbf{H}}$ and $\bar{\mathbf{B}}$ are the same. Thus $\bar{\mathbf{H}}\bar{\mathbf{H}}^T = \bar{\mathbf{B}}\bar{\mathbf{B}}^T$. In addition,

$$\mathcal{D}(\hat{\mathbf{V}},\mathscr{C}) = \|\hat{\mathbf{V}} - \hat{\mathbf{V}}\bar{\mathbf{H}}\bar{\mathbf{H}}^T\|_{\mathrm{F}}^2 = \|\mathbf{B}^T - \mathbf{B}^T\bar{\mathbf{B}}\bar{\mathbf{B}}^T\|_{\mathrm{F}}^2 \tag{22}$$
$$= \|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\bar{\mathbf{U}}_{N-K}^T\|_{\mathrm{F}}^2 = \|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\|_{\mathrm{F}}^2. \tag{23}$$

Note the eigendecompositions of $\bar{\mathbf{L}}$ and $\mathbf{L}$ can be written as

$$\bar{\mathbf{L}} = \bar{\mathbf{U}}_{N-K}\bar{\mathbf{\Sigma}}_{N-K}\bar{\mathbf{U}}_{N-K}^T + \bar{\mathbf{B}}\bar{\mathbf{\Lambda}}\bar{\mathbf{B}}^T, \tag{24}$$
$$\mathbf{L} = \mathbf{U}_{N-K}\mathbf{\Sigma}_{N-K}\mathbf{U}_{N-K}^T + \mathbf{B}\mathbf{\Lambda}\mathbf{B}^T. \tag{25}$$

By $|\lambda_{N-K+1}(\mathbf{L}) - \lambda_{N-K+1}(\bar{\mathbf{L}})| \leq \|\mathbf{E}\|_2$, we have

$$\lambda_{N-K+1}(\mathbf{L}) \leq \lambda_{N-K+1}(\bar{\mathbf{L}}) + \|\mathbf{E}\|_2. \tag{26}$$

Then if $\eta > 0$, we have that the eigenvalues of $\bar{\mathbf{\Sigma}}_{N-K}$ are contained in the interval $[N\min_k g_k, N\max_k g_k]$, and the eigenvalues of $\mathbf{\Lambda}$ are excluded from $[N\min_k g_k - \eta, N\max_k g_k + \eta]$. By the Davis-Kahan theorem,

$$\|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\|_2 \leq \frac{\|\mathbf{E}\|_2}{\eta}. \tag{27}$$

We further obtain that

$$\mathcal{D}(\hat{\mathbf{V}},\mathscr{C}) = \|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\|_{\mathrm{F}}^2 \tag{28}$$
$$\leq (K-1)\|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\|_2^2 \leq \frac{(K-1)\|\mathbf{E}\|_2^2}{\eta^2}, \tag{29}$$

where $\|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\|_{\mathrm{F}} \leq \sqrt{K-1}\|\mathbf{B}^T\bar{\mathbf{U}}_{N-K}\|_2$ because that $\mathbf{B}^T\bar{\mathbf{U}}_{N-K} \in \mathbb{R}^{K\times(N-K)}$ contains at most $K-1$ non-zero singular values (the last row of it contains all zeros). ■

4.4. Proof of the Main Theorem

Based on the above lemmas, we present our final proof.

Proof of Theorem 2 Let $\mathbf{B}_{K-1} = \mathbf{B}(:,1:(K-1))$. We have $\hat{\mathbf{V}} = [\mathbf{B}_{K-1}, \frac{\mathbf{e}}{\sqrt{N}}]^T$ and note that the sum of the entries of each column of $\mathbf{B}_{K-1}$ is 0. The centralized matrix of $\hat{\mathbf{V}}$ is

$$\hat{\mathbf{Z}} := [\mathbf{B}_{K-1}, \mathbf{0}]^T \in \mathbb{R}^{K\times N}. \tag{30}$$

For $\hat{\mathbf{S}} = \hat{\mathbf{Z}}^T\hat{\mathbf{Z}}$, we have $\lambda_{K-1}(\hat{\mathbf{S}}) = 1$ and $\lambda_n(\hat{\mathbf{S}}) = 0$ for $n \geq K$. Therefore, $\frac{\mathcal{D}(\hat{\mathbf{V}},\mathscr{C}) - \lambda_K(\hat{\mathbf{S}})}{\lambda_{K-1}(\hat{\mathbf{S}}) - \lambda_K(\hat{\mathbf{S}})} = \mathcal{D}(\hat{\mathbf{V}},\mathscr{C})$. Let $p_{\min} = \min_k \frac{n_k}{N}$. By the concentration inequality for sub-Gaussian random variables [16], for any $0 < \epsilon \leq \frac{w_{\min}}{2}$, there is a positive constant c depending on $\{w_k\}_{k\in[K]}$ such that

$$\mathbb{P}\left(\left|\frac{n_k}{N} - w_k\right| \geq \epsilon\right) \leq e\exp(-cN\epsilon^2) \text{ for } k \in [K]. \tag{31}$$

Note that $\zeta(\cdot)$ is monotonically increasing on $[0, \frac{K-1}{2}]$. Then if there exists $\epsilon > 0$ and $\xi \in (0,1)$, such that $\frac{4(K-1)T/\xi}{\left(\min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi} - \epsilon\right)^2} < \zeta(w_{\min} - \epsilon)$ and $\epsilon \leq \min\{\frac{w_{\min}}{2}, \min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi}\}$, we have that with probability at least $1 - \xi - Ke\exp\left(-cN\epsilon^2\right)$,

$$\mathcal{D}(\hat{\mathbf{V}},\mathscr{C}) \leq \frac{(K-1)\|\mathbf{E}\|_2^2}{\eta^2} \tag{32}$$
$$\leq \frac{4(K-1)T/\xi}{\left(\min_k c_k - \sqrt{2K}S_o - 2\sqrt{T}/\sqrt{\xi} - \epsilon\right)^2} \tag{33}$$
$$\leq \zeta(w_{\min} - \epsilon) \leq \zeta(p_{\min}). \tag{34}$$

Or equivalently, $\tau\left(\mathcal{D}(\hat{\mathbf{V}},\mathscr{C})\right) \leq p_{\min}$. By Lemma 1, we obtain the error bound given in (7). ■

4.5. Discussion about the Special Case

Considering the special case described in Remark 1. Let $S_d = s_{kk}$ for $k \in [K]$ and $S_o = s_{ij}$ for $i \neq j$. We have that for this special case, $\lambda_{N-K+1}(\bar{\mathbf{L}}) = \ldots = \lambda_{N-1}(\bar{\mathbf{L}}) = NS_o$ and the separability assumption (6) can be modified to be written as

$$\frac{4(K-1)T/\xi}{\left(w_{\min}(S_d - S_o) - 2\sqrt{T}/\sqrt{\xi}\right)^2} < \zeta(w_{\min}). \tag{35}$$

By Taylor approximation, we have that

$$S_d - S_o \approx \phi d^2, \quad \sqrt{T} \approx 2\sqrt{2F}\phi\sigma^2. \tag{36}$$

Note that $\zeta(w_{\min}) \geq \frac{w_{\min}}{2}$. Simplifying (35), we obtain (8).

5. EXTENSIONS

We may consider spectral clustering with normalized graph Laplacians. For example, the normalized graph Laplacian $\mathbf{L}_{\mathrm{rw}}$ is defined as

$$\mathbf{L}_{\mathrm{rw}} := \mathbf{D}^{-1}\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1}\mathbf{A}. \tag{37}$$

It is easy to derive similar theoretical results for such a normalized version of spectral clustering. In more detail, let $\bar{\mathbf{L}}_{\mathrm{rw}} := \bar{\mathbf{D}}^{-1}\bar{\mathbf{L}}$ and $\mathbf{E}_{\mathrm{rw}} = \mathbf{L}_{\mathrm{rw}} - \bar{\mathbf{L}}_{\mathrm{rw}}$. We make use of the inequality that $\|\bar{\mathbf{L}}_{\mathrm{rw}} - \mathbf{L}_{\mathrm{rw}}\|_2 \leq \|\bar{\mathbf{D}}^{-1} - \mathbf{D}^{-1}\|_2\|\bar{\mathbf{L}}\|_2 + \|\mathbf{D}^{-1}\|_2\|\bar{\mathbf{L}} - \mathbf{L}\|_2$ to derive an upper bound for $\|\mathbf{E}_{\mathrm{rw}}\|_2$, which is approximately $\frac{3\sqrt{T}/\sqrt{\xi}}{\min_k c_k - \sqrt{T}/\sqrt{\xi}}$, under certain conditions. We similarly analyze the spectral properties of $\bar{\mathbf{L}}_{\mathrm{rw}}$. Then we utilize the Davis-Kahan perturbation theorem and Lemma 1 to derive a theorem which is similar to Theorem 2. Due to space limit, we omit the details.

6. REFERENCES

[1] S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 129–137, 1982.

[2] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *NIPS*, 2002, pp. 849–856.

[3] F. McSherry, "Spectral partitioning of random graphs," in *FOCS*, 2001.

[4] K. Rohe, S. Chatterjee, and B. Yu, "Spectral clustering and the high-dimensional stochastic blockmodel," *Ann. Stat.*, vol. 39, no. 4, pp. 1878–1915, 2011.

[5] L. Huang, D. Yan, N. Taft, and M. I. Jordan, "Spectral clustering with perturbed data," in *NIPS*, 2009, pp. 705–712.

[6] S. Balakrishnan, M. Xu, A. Krishnamurthy, and A. Singh, "Noise thresholds for spectral clustering," in *NIPS*, 2011, pp. 954–962.

[7] C. Ding, X. He, and H. D. Simon, "On the equivalence of nonnegative matrix factorization and spectral clustering," in *SDM*, 2005, pp. 606–610.

[8] J. Lei and A. Rinaldo, "Consistency of spectral clustering in stochastic block models," *Ann. Stat.*, vol. 43, no. 1, pp. 215–237, 2015.

[9] U. Von Luxburg, M. Belkin, and O. Bousquet, "Consistency of spectral clustering," *Ann. Stat.*, pp. 555–586, 2008.

[10] U. Von Luxburg, "A tutorial on spectral clustering," *Stat. Comput.*, vol. 17, no. 4, pp. 395–416, 2007.

[11] C. Ding and X. He, "k-means clustering via principal component analysis," in *ICML*, 2004.

[12] M. Meilă, "Comparing clusterings: An axiomatic view," in *ICML*, 2005, pp. 577–584.

[13] M. Meilă, "The uniqueness of a good optimum for k-means," in *ICML*, 2006, pp. 625–632.

[14] Z. Liu and V. Y. F. Tan, "The informativeness of k-means and dimensionality reduction for learning mixture models," *arXiv preprint arXiv:1703.10534*, 2017.

[15] C. Davis and W. M. Kahan, "The rotation of eigenvectors by a perturbation. III," *SIAM J. Numer. Anal.*, vol. 7, no. 1, pp. 1–46, 1970.

[16] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," *arXiv preprint arXiv:1011.3027*, 2010.