
Spectral Clustering via the Power Method - Provably

Christos Boutsidis

Yahoo, 229 West 43rd Street, New York, NY, USA.

BOUTSIDIS@YAHOO-INC.COM

Alex Gittens

International Computer Science Institute, Berkeley, CA, USA.

GITTENS@ICSI.BERKELEY.EDU

Prabhanjan Kambadur

Bloomberg L.P., 731 Lexington Avenue, New York, 10022, USA.

PKAMBADUR@BLOOMBERG.NET

Abstract

Spectral clustering is one of the most important algorithms in data mining and machine intelligence; however, its computational complexity limits its application to truly large scale data analysis. The computational bottleneck in spectral clustering is computing a few of the top eigenvectors of the (normalized) Laplacian matrix corresponding to the graph representing the data to be clustered. One way to speed up the computation of these eigenvectors is to use the “power method” from the numerical linear algebra literature. Although the power method has been empirically used to speed up spectral clustering, the theory behind this approach, to the best of our knowledge, remains unexplored. This paper provides the *first* such rigorous theoretical justification, arguing that a small number of power iterations suffices to obtain near-optimal partitionings using the approximate eigenvectors. Specifically, we prove that solving the k -means clustering problem on the approximate eigenvectors obtained via the power method gives an additive-error approximation to solving the k -means problem on the optimal eigenvectors.

1. Introduction

Consider clustering the points in Figure 1. The data in this space are non-separable and there is no apparent clustering metric which can be used to recover this clustering structure. In particular, the two clusters have the same centers (centroids); hence, distance-based clustering methods such as k -means (Ostrovsky et al., 2006) will fail. Motivated by

Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 2015. JMLR: W&CP volume 37. Copyright 2015 by the author(s).




Figure 1. 2-D data amenable to spectral clustering.

such shortcomings of traditional clustering approaches, researchers have produced a body of more flexible and data-adaptive clustering approaches, now known under the umbrella of *spectral clustering*. The crux of these approaches is to model the points to be clustered as vertices of a graph, where weights on edges connecting the vertices are assigned according to some similarity measure between the points. Next, a new, hopefully separable, representation of the points is formed by using the eigenvectors of the (normalized) Laplacian matrix associated with this similarity graph. This new, typically low-dimensional, representation is often called “spectral embedding” of the points. We refer the reader to (Fiedler, 1973; Von Luxburg, 2007; Shi & Malik, 2000b) for the foundations of spectral clustering and to (Belkin & Niyogi, 2001; Ng et al., 2002; Liu & Zhang, 2004; Zelnik-Manor & Perona, 2004; Smyth & White, 2005) for applications in data mining and machine learning. We explain spectral clustering and the baseline algorithm in detail in Section 2.1.

The computational bottleneck in spectral clustering is the computation of the eigenvectors of the Laplacian matrix. Motivated by the need for faster algorithms to compute these eigenvectors, several techniques have been developed in order to speedup this computation (Spielman & Teng, 2009; Yan et al., 2009; Fowlkes et al., 2004; Pavan & Pelillo, 2005; Bezdek et al., 2006; Wang et al., 2009; Nyström, 1930; Baker, 1977). Perhaps the most popular of the above mentioned techniques is the “power method” (Lin & Cohen, 2010). The convergence of the power method is theoretically well understood when it comes to measure the principal angle between the space spanned by the true and the approximate eigenvectors (see Theorem 8.2.4 in (Golub & Van Loan, 2012)). We refer readers to (Woodruff, 2014) for a rigorous theoretical analysis of the use of the power method for the low-rank matrix approximation problem. However, these results do not imply that the approximate eigenvectors of the power method are useful for spectral clustering.

Contributions. In this paper, we argue that the eigenvectors computed via the power method are useful for spectral clustering, and that the loss in clustering accuracy is small. We prove that solving the k -means problem on the approximate eigenvectors obtained via the power method gives an additive-error approximation to solving the k -means problem on the optimal eigenvectors (see Lemma 5 and Thm 6).

2. Background

2.1. Spectral Clustering

We first review one mathematical formulation of spectral clustering. Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$ be n points in d dimensions. The goal of clustering is to partition these points into k disjoint sets, for some given k . To this end, define a weighted undirected graph $G(V, E)$ with $|V| = n$ nodes and $|E|$ edges: each node in G corresponds to an \mathbf{x}_i ; the weight of each edge encodes the similarity between its end points. Let $\mathbf{W} \in \mathbb{R}^{n \times n}$ be the similarity matrix — $\mathbf{W}_{ij} = e^{-(\|\mathbf{x}_i - \mathbf{x}_j\|^2)/\sigma}$, $i \neq j$ and $\mathbf{W}_{ii} = 0$ — that gives the similarity between \mathbf{x}_i and \mathbf{x}_j . Here, σ is a tuning parameter. Given this setup, spectral clustering for $k = 2$ corresponds to the following graph partitioning problem:

Definition 1 (The Spectral Clustering Problem for $k = 2$ (Shi & Malik, 2000b)). *Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$ and $k = 2$ be given. Construct graph $G(V, E)$ as described in the text above. Find subgraphs A and B of G that minimize the following:*

$$\text{Ncut}(A, B) = \text{cut}(A, B) \cdot \left(\frac{1}{\text{assoc}(A, V)} + \frac{1}{\text{assoc}(B, V)} \right),$$

where, $\text{cut}(A, B) = \sum_{\mathbf{x}_i \in A, \mathbf{x}_j \in B} \mathbf{W}_{ij}$; $\text{assoc}(A, V) = \sum_{\mathbf{x}_i \in A, \mathbf{x}_j \in V} \mathbf{W}_{ij}$; $\text{assoc}(B, V) = \sum_{\mathbf{x}_i \in B, \mathbf{x}_j \in V} \mathbf{W}_{ij}$.

This definition generalizes to any $k > 2$ in a straightforward manner (we omit the details). Minimizing $\text{Ncut}(A, B)$ in a weighted undirected graph is an NP-Complete problem (see appendix in (Shi & Malik, 2000b) for proof). Motivated by this hardness result, Shi and Malik (Shi & Malik, 2000b) suggested a relaxation to this problem that is tractable in polynomial time using the Singular Value Decomposition (SVD). First, (Shi & Malik, 2000b) shows that for any G, A, B and partition vector $\mathbf{y} \in \mathbb{R}^n$ with $+1$ to the entries corresponding to A and -1 to the entries corresponding to B the following identity holds: $4 \cdot \text{Ncut}(A, B) = \mathbf{y}^T (\mathbf{D} - \mathbf{W}) \mathbf{y} / (\mathbf{y}^T \mathbf{D} \mathbf{y})$. Here, $\mathbf{D} \in \mathbb{R}^{n \times n}$ is the diagonal matrix of degree nodes: $\mathbf{D}_{ii} = \sum_j \mathbf{W}_{ij}$. Hence, the spectral clustering problem in Definition 1 can be restated as finding such an optimum partition vector \mathbf{y} , which, as we mentioned above, is an intractable problem. The real relaxation for spectral clustering asks for a real-valued vector $\mathbf{y} \in \mathbb{R}^n$:

Definition 2 (The real relaxation for the spectral clustering problem for $k = 2$ (Shi & Malik, 2000b)). *Given graph G with n nodes, adjacency matrix \mathbf{W} , and degrees matrix \mathbf{D} find $\mathbf{y} \in \mathbb{R}^n$ such that:*

$$\mathbf{y} = \underset{\mathbf{y} \in \mathbb{R}^n, \mathbf{y}^T \mathbf{D} \mathbf{1}_n}{\text{argmin}} \frac{(\mathbf{y}^T (\mathbf{D} - \mathbf{W}) \mathbf{y})}{(\mathbf{y}^T \mathbf{D} \mathbf{y})}.$$

Once such a \mathbf{y} is found, one can partition the graph into two subgraphs by looking at the signs of the elements in \mathbf{y} . When $k > 2$, one can compute k eigenvectors and then apply k -means clustering on the rows of a matrix, denoted as \mathbf{Y} , containing those eigenvectors in its columns.

Motivated by these observations, Ng et. al (Ng et al., 2002) (see also (Weiss, 1999)) suggested the following algorithm for spectral clustering¹ (inputs to the algorithm are the points $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$ and the number of clusters k).

1. Construct the similarity matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ as $\mathbf{W}_{ij} = e^{-(\|\mathbf{x}_i - \mathbf{x}_j\|^2)/\sigma}$ (for $i \neq j$); $\mathbf{W}_{ii} = 0$ and σ is given.
2. Construct $\mathbf{D} \in \mathbb{R}^{n \times n}$ as the diagonal matrix of degrees of the nodes: $\mathbf{D}_{ii} = \sum_j \mathbf{W}_{ij}$.
3. Construct $\tilde{\mathbf{W}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}^{-\frac{1}{2}} \in \mathbb{R}^{n \times n}$.²
4. Find the largest k eigenvectors of $\tilde{\mathbf{W}}$ and assign them as columns to a matrix $\mathbf{Y} \in \mathbb{R}^{n \times k}$.³
5. Apply k -means clustering on the rows of \mathbf{Y} , and use this clustering to cluster the original points accordingly.

¹Precisely, Ng et. al suggested an additional normalization step on \mathbf{Y} before applying k -means, i.e., normalize \mathbf{Y} to unit row norms, but we ignore this step for simplicity.

²Here, $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is the Laplacian matrix of G and $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{W}}$ is the so called normalized Laplacian matrix.

³The top k eigenvectors of $\mathbf{D}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}^{-\frac{1}{2}}$ correspond to the bottom k eigenvectors of $\mathbf{I}_n - \mathbf{D}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}^{-\frac{1}{2}}$.

This algorithm serves as our baseline for an “exact spectral clustering algorithm”. One way to speedup⁴ this baseline algorithm is to use the power method (Lin & Cohen, 2010) in Step 4 to quickly approximate the eigenvectors in \mathbf{Y} ; that is,

- **Power method:** Initialize $\mathbf{S} \in \mathbb{R}^{n \times k}$ with *i.i.d* random Gaussian variables. Let $\tilde{\mathbf{Y}} \in \mathbb{R}^{n \times k}$ contain the left singular vectors of the matrix

$$\mathbf{B} = (\tilde{\mathbf{W}}\tilde{\mathbf{W}}^T)^p \tilde{\mathbf{W}}\mathbf{S} = \tilde{\mathbf{W}}^{2p+1}\mathbf{S},$$

for some integer $p \geq 0$. Now, use $\tilde{\mathbf{Y}}$ instead of \mathbf{Y} in step 5 above.

The use of the power method to speedup eigenvector computation is not new. Power method is a classical technique in the numerical linear algebra literature (see Section 8.2.4 in (Golub & Van Loan, 2012)). Existing theoretical analysis provides sharp bounds for the error $\|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_2$ (see Theorem 8.2.4 in (Golub & Van Loan, 2012); this theorem assumes that \mathbf{S} has orthonormal columns and it is not perpendicular to \mathbf{Y} .) (Halko et al., 2011; Woodruff, 2014) also used the power method with random Gaussian initialization and applied it to the low-rank matrix approximation problem. The approximation bounds proved in those papers are for the term $\|\tilde{\mathbf{W}} - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\tilde{\mathbf{W}}\|_2$. To the best of our knowledge, none of these results indicates that the approximate eigenvectors are useful for spectral clustering purposes.

2.2. Connection to k -means

The previous algorithm indicates that spectral clustering turns out to be a k -means clustering problem on the rows of \mathbf{Y} , the matrix containing the bottom eigenvectors of the normalized Laplacian matrix. The main result of our paper is to prove that solving the k -means problem on $\tilde{\mathbf{Y}}$ and using this to cluster the rows in \mathbf{Y} gives a clustering which is as good as the clustering by solving the k -means problem on \mathbf{Y} . To this end, we need some background on k -means clustering; we present a linear algebraic view below.

For $i = 1 : n$, let $\mathbf{y}_i \in \mathbb{R}^k$ be a row of \mathbf{Y} as a column vector. Hence,

$$\mathbf{Y} = \begin{bmatrix} -\mathbf{y}_1^T & - \\ -\mathbf{y}_2^T & - \\ \vdots & \\ -\mathbf{y}_n^T & - \end{bmatrix} \in \mathbb{R}^{n \times k}.$$

⁴This can be implemented in $O(n^2kp + k^2n)$ time, as we need $O(n^2kp)$ time to implement all the matrix-matrix multiplications (right-to-left) and another $O(k^2n)$ time to find $\tilde{\mathbf{Y}}$. As we discuss below, selecting $p \approx O(\ln(kn))$ and assuming that the multiplicative spectral gap of between the k^{th} and $(k+1)^{\text{th}}$ eigenvalue of $\tilde{\mathbf{W}}$ is large suffices to get very accurate clusterings. This leads to an $O(kn^2 \ln(kn))$ runtime for this step.

Let k be the number of clusters. One can define a partition of the rows of \mathbf{Y} by a cluster indicator matrix $\mathbf{X} \in \mathbb{R}^{n \times k}$. Each column $j = 1, \dots, k$ of \mathbf{X} represents a cluster. Each row $i = 1, \dots, n$ indicates the cluster membership of \mathbf{y}_i . So, $\mathbf{X}_{ij} = 1/\sqrt{s_j}$, if and only if the data point \mathbf{y}_i is in the j th cluster ($s_j = \|\mathbf{X}^{(j)}\|_0$; $\mathbf{X}^{(j)}$ is the j th column of \mathbf{X} and $\|\mathbf{X}^{(j)}\|_0$ denotes the number of non-zero elements of $\mathbf{X}^{(j)}$). We formally define the k -means problem as follows:

Definition 3. [THE k -MEANS CLUSTERING PROBLEM] Given $\mathbf{Y} \in \mathbb{R}^{n \times k}$ (representing n data points – rows – described with respect to k features – columns) and a positive integer k denoting the number of clusters, find the indicator matrix $\mathbf{X}_{\text{opt}} \in \mathbb{R}^{n \times k}$ which satisfies,

$$\mathbf{X}_{\text{opt}} = \underset{\mathbf{X} \in \mathcal{X}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{X}\mathbf{X}^T\mathbf{Y}\|_{\text{F}}^2.$$

Here, \mathcal{X} denotes the set of all $m \times k$ indicator matrices \mathbf{X} . Also, we will denote

$$\|\mathbf{Y} - \mathbf{X}_{\text{opt}}\mathbf{X}_{\text{opt}}^T\mathbf{Y}\|_{\text{F}}^2 := \text{F}_{\text{opt}}.$$

This definition is equivalent to the more traditional definition involving sum of squared distances of points from cluster centers (Boutsidis & Magdon-Ismail, 2013; Cohen et al., 2014) (we omit the details). Next, we formalize the notion of a “ k -means approximation algorithm”.

Definition 4. [k -MEANS APPROXIMATION ALGORITHM] An algorithm is called a “ γ -approximation” for the k -means clustering problem ($\gamma \geq 1$) if it takes inputs the dataset $\mathbf{Y} \in \mathbb{R}^{n \times k}$ and the number of clusters k , and returns an indicator matrix $\mathbf{X}_\gamma \in \mathbb{R}^{n \times k}$ such that w.p. $1 - \delta_\gamma$,

$$\|\mathbf{Y} - \mathbf{X}_\gamma\mathbf{X}_\gamma^T\mathbf{Y}\|_{\text{F}}^2 \leq \gamma \min_{\mathbf{X} \in \mathcal{X}} \|\mathbf{Y} - \mathbf{X}\mathbf{X}^T\mathbf{Y}\|_{\text{F}}^2 = \gamma \cdot \text{F}_{\text{opt}}.$$

An example of such an approximation algorithm is in (Kumar et al., 2004) with $\gamma = 1 + \varepsilon$ ($0 < \varepsilon < 1$) and δ_γ a constant in $(0, 1)$. The corresponding running time is $O(nk \cdot 2^{(k/\varepsilon)^{O(1)}})$. A trivial algorithm with $\gamma = 1$ but running time $\Omega(n^k)$ is to try all possible k -clusterings of the rows of \mathbf{Y} and keep the best.

3. Main result

Next, we argue that applying a k -means approximation algorithm on \mathbf{Y} and $\tilde{\mathbf{Y}}$ gives approximately the same clustering results for a sufficiently large number of power iterations p . Hence, the eigenvectors found by the power method do not sacrifice the accuracy of the exact spectral clustering algorithm. This is formally shown in Theorem 6. However, the key notion of approximation between the exact and the approximate eigenvectors is captured in Lemma 5:

Lemma 5. [See Section 4.1 for proof] For any $\varepsilon, \delta > 0$, let

$$p \geq \frac{\frac{1}{2} \ln \left(4 \cdot n \cdot \varepsilon^{-1} \cdot \delta^{-1} \cdot \sqrt{k} \right)}{\ln(\gamma_k)},$$

where

$$\gamma_k = \frac{\sigma_k(\tilde{\mathbf{W}})}{\sigma_{k+1}(\tilde{\mathbf{W}})},$$

is the multiplicative eigen-gap between the k -th and the $(k+1)$ -th singular value of $\tilde{\mathbf{W}}$. Then, with probability at least $1 - e^{-2n} - 2.35\delta$: $\|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_{\mathbb{F}}^2 \leq \varepsilon^2$.

In words, for a sufficiently large value of p , the orthogonal projection operators on $\text{span}(\mathbf{Y})$ and $\text{span}(\tilde{\mathbf{Y}})$ are bounded, in Frobenius norm, for an arbitrarily small ε . Here and throughout the paper we reserve $\sigma_1(\tilde{\mathbf{W}})$ to denote the largest singular value of $\tilde{\mathbf{W}}$:

$$\sigma_1(\tilde{\mathbf{W}}) \geq \sigma_2(\tilde{\mathbf{W}}) \geq \dots \geq \sigma_n(\tilde{\mathbf{W}}) \geq 0;$$

and similarly for $\tilde{\mathbf{L}}$:

$$\lambda_1(\tilde{\mathbf{L}}) \geq \lambda_2(\tilde{\mathbf{L}}) \geq \dots \geq \lambda_n(\tilde{\mathbf{L}}).$$

Next, we present our main theorem (see Section 4.2 for the proof).

Theorem 6. Construct $\tilde{\mathbf{Y}}$ via the power method with

$$p \geq \frac{\frac{1}{2} \cdot \ln \left(4 \cdot n \cdot \varepsilon^{-1} \cdot \delta^{-1} \cdot \sqrt{k} \right)}{\ln(\gamma_k)},$$

where

$$\gamma_k = \frac{\sigma_k(\tilde{\mathbf{W}})}{\sigma_{k+1}(\tilde{\mathbf{W}})} = \frac{1 - \sigma_{n-k+1}(\tilde{\mathbf{L}})}{1 - \sigma_{n-k}(\tilde{\mathbf{L}})}.$$

Here, $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{W}}$ is the normalized Laplacian matrix. Consider running on the rows of $\tilde{\mathbf{Y}}$ a γ -approximation k -means algorithm with failure probability δ_γ . Let the outcome be a clustering indicator matrix $\mathbf{X}_{\tilde{\gamma}} \in \mathbb{R}^{n \times k}$. Also, let \mathbf{X}_{opt} be the optimal clustering indicator matrix for \mathbf{Y} . Then, with probability at least $1 - e^{-2n} - 2.35\delta - \delta_\gamma$,

$$\|\mathbf{Y} - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \mathbf{Y}\|_{\mathbb{F}}^2 \leq (1+4\varepsilon) \cdot \gamma \cdot \|\mathbf{Y} - \mathbf{X}_{\text{opt}} \mathbf{X}_{\text{opt}}^T \mathbf{Y}\|_{\mathbb{F}}^2 + 4\varepsilon^2.$$

3.1. Discussion

Several remarks are necessary regarding our main theorem. First of all, notice that the notion of approximation in clustering quality is with respect to the objective value of k -means; indeed, this is often the case in approximation algorithms for k -means clustering (Cohen et al., 2014). We

acknowledge that it would have been better to obtain results on how well $\mathbf{X}_{\tilde{\gamma}}$ approximates \mathbf{X}_{opt} directly, for example, via bounding the error $\|\mathbf{X}_{\tilde{\gamma}} - \mathbf{X}_{\text{opt}}\|_{\mathbb{F}}^2$; however, such results are notoriously difficult to obtain since this is a combinatorial objective.

Next, let us take a more careful look at the effect of the parameter γ_k . It suffices to discuss this effect for two cases: 1) $\gamma_k = 1$; and 2) $\gamma_k > 1$. To this end, we need to use a relation between the eigenvalues of $\tilde{\mathbf{W}}$ and $\tilde{\mathbf{L}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{L} \mathbf{D}^{-\frac{1}{2}}$ (see proof of theorem for explanation): for all $i = 1, 2, \dots, n$, the relation is:

$$\sigma_i(\tilde{\mathbf{W}}) = 1 - \sigma_{n-i+1}(\tilde{\mathbf{L}}).$$

3.1.1. $\gamma_k = 1$

First, we argue that the case $\gamma_k = 1$ is not interesting from a spectral clustering perspective and hence, we can safely assume that this will not occur in practical scenarios. γ_k is the multiplicative eigen-gap between the k th and the $(k+1)$ th eigenvalue of $\tilde{\mathbf{W}}$:

$$\gamma_k = \frac{\sigma_k(\tilde{\mathbf{W}})}{\sigma_{k+1}(\tilde{\mathbf{W}})} = \frac{1 - \lambda_{n-k+1}(\tilde{\mathbf{L}})}{1 - \lambda_{n-k}(\tilde{\mathbf{L}})}.$$

The folklore belief (see end of Section 4.3 in (Lee et al., 2012)) in spectral clustering says that a good k to select in order to cluster the data via spectral clustering is when

$$\lambda_{n-k+1}(\tilde{\mathbf{L}}) \ll \lambda_{n-k}(\tilde{\mathbf{L}}).$$

In this case, the spectral gap is sufficiently large and γ_k will not be close to one; hence, a small number of power iterations p will be enough to obtain accurate results. To summarize, if $\gamma_k = 1$, it does not make sense to perform spectral clustering with k clusters and the user should look for a $k' > k$ such that the gap in the spectrum, $\gamma_{k'}$ is not approaching and at the very least is strictly larger than 1.

This gap assumption is not surprising from a linear algebraic perspective as well. It is well known that in order the power iteration to succeed to find the eigenvectors of any symmetric matrix, the multiplicative eigen-gap in the spectrum should be sufficiently large, as otherwise the power method will not be able to distinguish the k th eigenvector from the $(k+1)$ th eigenvector. For a more detailed discussion of this we refer the reader to Thm 8.2.4 in (Golub & Van Loan, 2012).

3.1.2. $\gamma_k > 1$

We remark that the graph we construct to pursue spectral clustering is a complete graph, hence it is connected; a basic fact in spectral graph theory (see, for example, (Lee

et al., 2012)) says that the smallest eigenvalue of the Laplacian matrix \mathbf{L} of the graph equals to zero ($\lambda_n(\mathbf{L}) = 0$) if and only if the graph is connected. An extension of this fact says that the number of disconnected components in the graph equals the multiplicity of the eigenvalue zero in the Laplacian matrix \mathbf{L} . When this happens, we have

$$\sigma_n(\mathbf{L}) = \sigma_{n-1}(\mathbf{L}) = \dots = \sigma_{n-k+1}(\mathbf{L}) = 0,$$

and, correspondingly,

$$\sigma_n(\tilde{\mathbf{L}}) = \sigma_{n-1}(\tilde{\mathbf{L}}) = \dots = \sigma_{n-k+1}(\tilde{\mathbf{L}}) = 0.$$

What happens, however, when the k smallest eigenvalues of the normalized Laplacian are not zero but close to zero? Cheeger’s inequality and extensions in (Lee et al., 2012) indicate that as those eigenvalues approaching zero, the graph is approaching a graph with k disconnected components, hence clustering such graphs should be “easy”, that is, the number of power iterations p should be small. We formally argue about this statement below. First, we state a version of the Cheeger’s inequality that explains the situation for the $k = 2$ case (we omit the details of the high order extensions in (Lee et al., 2012)). The facts below can be found in Section 1.2 in (Bandeira et al., 2013)). Recall also the graph partitioning problem in Definition 1. Let the minimum value for $\text{Ncut}(A, B)$ be obtained from a partition into two sets A_{opt} and B_{opt} , then, implies

$$\frac{1}{2}\lambda_{n-1}(\tilde{\mathbf{L}}) \leq \text{Ncut}(A_{opt}, B_{opt}) \leq 2\sqrt{2\lambda_{n-1}(\tilde{\mathbf{L}})}.$$

Hence, if $\lambda_{n-1}(\tilde{\mathbf{L}}) \rightarrow 0$, then $\text{Ncut}(A_{opt}, B_{opt}) \rightarrow 0$, which makes the clustering problem “easy”, hence amenable to a small number of power iterations.

Now, we formally derive the relation of p as a function of $\sigma_{n-k+1}(\tilde{\mathbf{L}})$. Towards this end, fix all values ($n, \varepsilon, \delta, \sigma_{n-k}(\tilde{\mathbf{L}})$) to constants, e.g., $n = 10^3$, $\varepsilon = 10^{-3}$, $\delta = 10^{-2}$, and $\sigma_{n-k}(\tilde{\mathbf{L}}) = 1/2$. Also, $x = \sigma_{n-k+1}(\tilde{\mathbf{L}})$. Then,

$$p := f(x) = \frac{1}{2} \frac{\ln(4 \cdot 10^9)}{\ln(2 - 2 \cdot x)}.$$

Simple calculus arguments show that as $0 \leq x < 1/2$ (this range for x is required to make sure that $\gamma_k > 1$) increases, then $f(x)$ also increases, which confirms the expectation that for an eigenvalue $x := \sigma_{n-k+1}(\tilde{\mathbf{L}})$ approaching zero, the number of power iterations is approaching zero as well. We plot $f(x)$ in Figure 2. The number of power iterations is always small since the dependence of p on $\sigma_{n-k+1}(\tilde{\mathbf{L}})$ is logarithmic.

3.1.3. DEPENDENCE ON ε

Finally, we remark that the approximation bound in the theorem indicates that the loss in clustering accuracy can be




Figure 2. Number of power iterations p vs the eigenvalue $x := \sigma_{n-k+1}(\tilde{\mathbf{L}})$ of the normalized Laplacian.

made arbitrarily small since the dependence on ε^{-1} is logarithmic with respect to the number of power iterations. In particular, when $\varepsilon \leq \|\mathbf{Y} - \mathbf{X}_{opt}\mathbf{X}_{opt}^T\mathbf{Y}\|_F^2$, we have a relative error bound: $\|\mathbf{Y} - \mathbf{X}_{\tilde{\gamma}}\mathbf{X}_{\tilde{\gamma}}^T\mathbf{Y}\|_F^2 \leq$

$$\begin{aligned} &\leq ((1 + 4\varepsilon) \cdot \gamma + 4\varepsilon) \|\mathbf{Y} - \mathbf{X}_{opt}\mathbf{X}_{opt}^T\mathbf{Y}\|_F^2 \\ &\leq (1 + 8\varepsilon) \cdot \gamma \cdot \|\mathbf{Y} - \mathbf{X}_{opt}\mathbf{X}_{opt}^T\mathbf{Y}\|_F^2, \end{aligned}$$

where the last relation uses $1 \leq \gamma$. One might wonder that to achieve this relative error performance, $\varepsilon \rightarrow 0$, in which case it appears that p is very large. However, this is not true since the event that $\varepsilon \rightarrow 0$ is necessary occurs only when $\|\mathbf{Y} - \mathbf{X}_{opt}\mathbf{X}_{opt}^T\mathbf{Y}\|_F^2 \rightarrow 0$, which happens when $\mathbf{Y} \approx \mathbf{X}_{opt}$. In this case, from the discussion in the previous section, we have $\sigma_{n-k+1}(\tilde{\mathbf{L}}) \rightarrow 0$, because \mathbf{Y} is “close” to an indicator matrix if and only if the graph is “close” to having k disconnected components, which itself happens if and only if $\sigma_{n-k+1}(\tilde{\mathbf{L}}) \rightarrow 0$. It is easy now to calculate that, when this happens, $p \rightarrow 0$ because, from standard calculus arguments, we can derive:

$$\lim_{x \rightarrow 0} \frac{\frac{1}{2} \cdot \ln\left(4 \cdot n \cdot \frac{1}{x} \cdot \delta^{-1} \cdot \sqrt{k}\right)}{\ln\left(\frac{1-x}{1-\sigma_{n-k}(\tilde{\mathbf{L}})}\right)} = 0.$$

4. Proofs

4.1. Proof of Lemma 5

4.1.1. PRELIMINARIES

We first introduce the notation that we use throughout the paper. $\mathbf{A}, \mathbf{B}, \dots$ are matrices; \mathbf{a}, \dots are column vectors. \mathbf{I}_n is the $n \times n$ identity matrix; $\mathbf{0}_{m \times n}$ is the $m \times n$ matrix of zeros; $\mathbf{1}_n$ is the $n \times 1$ vector of ones. The Frobenius and the spectral matrix-norms are $\|\mathbf{A}\|_F^2 = \sum_{i,j} \mathbf{A}_{ij}^2$ and $\|\mathbf{A}\|_2 = \max_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$, respectively. The thin (compact) SVD of $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ is

$$\mathbf{A} = \underbrace{\begin{pmatrix} \mathbf{U}_k & \mathbf{U}_{\rho-k} \end{pmatrix}}_{\mathbf{U}_A \in \mathbb{R}^{m \times \rho}} \underbrace{\begin{pmatrix} \Sigma_k & \mathbf{0} \\ \mathbf{0} & \Sigma_{\rho-k} \end{pmatrix}}_{\Sigma_A \in \mathbb{R}^{\rho \times \rho}} \underbrace{\begin{pmatrix} \mathbf{V}_k^T \\ \mathbf{V}_{\rho-k}^T \end{pmatrix}}_{\mathbf{V}_A^T \in \mathbb{R}^{\rho \times n}}, \quad (1)$$

with singular values $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_k(\mathbf{A}) \geq \sigma_{k+1}(\mathbf{A}) \geq \dots \geq \sigma_\rho(\mathbf{A}) > 0$. The matrices $\mathbf{U}_k \in \mathbb{R}^{m \times k}$ and $\mathbf{U}_{\rho-k} \in \mathbb{R}^{m \times (\rho-k)}$ contain the left singular vectors of \mathbf{A} ; and, similarly, the matrices $\mathbf{V}_k \in \mathbb{R}^{n \times k}$ and $\mathbf{V}_{\rho-k} \in \mathbb{R}^{n \times (\rho-k)}$ contain the right singular vectors. $\Sigma_k \in \mathbb{R}^{k \times k}$ and $\Sigma_{\rho-k} \in \mathbb{R}^{(\rho-k) \times (\rho-k)}$ contain the singular values of \mathbf{A} . Also, $\mathbf{A}^\dagger = \mathbf{V}_A \Sigma_A^{-1} \mathbf{U}_A^T$ denotes the pseudo-inverse of \mathbf{A} . For a symmetric positive definite matrix (SPSD) $\mathbf{A} = \mathbf{B}\mathbf{B}^T$, $\lambda_i(\mathbf{A}) = \sigma_i^2(\mathbf{B})$ denotes the i -th eigenvalue of \mathbf{A} .

4.1.2. INTERMEDIATE LEMMAS

To prove Lemma 5, we need the following simple fact.

Lemma 7. For any $\mathbf{Y} \in \mathbb{R}^{n \times k}$ and $\tilde{\mathbf{Y}} \in \mathbb{R}^{n \times k}$ with $\mathbf{Y}^T \mathbf{Y} = \tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}} = \mathbf{I}_k$:

$$\|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F^2 \leq 2k \|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_2^2.$$

Proof. This is because $\|\mathbf{X}\|_F^2 \leq \text{rank}(\mathbf{X}) \|\mathbf{X}\|_2^2$, for any \mathbf{X} , and the fact that $\text{rank}(\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T) \leq 2k$. To justify this, notice that $\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T$ can be written as the product of two matrices each with rank at most $2k$:

$$\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T = \begin{pmatrix} \mathbf{Y} & \tilde{\mathbf{Y}} \end{pmatrix} \begin{pmatrix} \mathbf{Y}^T \\ -\tilde{\mathbf{Y}}^T \end{pmatrix}. \quad \blacksquare$$

We also need the following result, which appeared as Lemma 7 in (Boutsidis & Magdon-Ismail, 2014).

Lemma 8. For any matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ with rank at least k , let $p \geq 0$ be an integer and draw $\mathbf{S} \in \mathbb{R}^{n \times k}$, a matrix of i.i.d. standard Gaussian random variables. Fix $\delta \in (0, 1)$ and $\varepsilon \in (0, 1)$. Let $\Omega_1 = (\mathbf{A}\mathbf{A}^T)^p \mathbf{A}\mathbf{S}$, and $\Omega_2 = \mathbf{A}_k$. If

$$p \geq \ln(4n\varepsilon^{-1}\delta^{-1}) \ln^{-1} \left(\frac{\sigma_k(\mathbf{A})}{\sigma_{k+1}(\mathbf{A})} \right),$$

then w.p. $1 - e^{-2n} - 2.35\delta$, $\|\Omega_1 \Omega_1^\dagger - \Omega_2 \Omega_2^\dagger\|_2^2 \leq \varepsilon^2$.

4.1.3. CONCLUDING THE PROOF OF LEMMA 5

Applying Lemma 8 with $\mathbf{A} = \tilde{\mathbf{W}}$ and

$$p \geq \ln(4n\varepsilon^{-1}\delta^{-1}) \ln^{-1} \left(\frac{\sigma_k(\tilde{\mathbf{W}})}{\sigma_{k+1}(\tilde{\mathbf{W}})} \right),$$

gives

$$\|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F^2 \leq \varepsilon^2.$$

Rescaling $\varepsilon' = \varepsilon/\sqrt{k}$, i.e., choosing p as

$$p \geq \ln(4n\varepsilon^{-1}\delta^{-1}\sqrt{k}) \ln^{-1} \left(\frac{\sigma_k(\tilde{\mathbf{W}})}{\sigma_{k+1}(\tilde{\mathbf{W}})} \right),$$

gives

$$\|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F^2 \leq \varepsilon^2/k.$$

Combining this bound with Lemma 7 we obtain:

$$\|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F^2 \leq 2k \|\mathbf{Y}\mathbf{Y}^T - \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_2^2 \leq 2k \cdot (\varepsilon^2/k) = \varepsilon^2,$$

as advertised.

4.2. Proof of Theorem 6

Let

$$\mathbf{Y}\mathbf{Y}^T = \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T + \mathbf{E},$$

where \mathbf{E} is an $n \times n$ matrix with $\|\mathbf{E}\|_F \leq \varepsilon$ (this follows after taking square root on both sides in the inequality of Lemma 5). Next, we manipulate the term $\|\mathbf{Y} - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \mathbf{Y}\|_F$ as follows

$$\begin{aligned} \|\mathbf{Y} - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \mathbf{Y}\|_F &= \\ &= \|\mathbf{Y}\mathbf{Y}^T - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \mathbf{Y}\mathbf{Y}^T\|_F \\ &= \|\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T + \mathbf{E} - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T (\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T + \mathbf{E})\|_F \\ &\leq \|\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F + \|(\mathbf{I}_n - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T) \mathbf{E}\|_F \\ &\leq \|\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F + \|\mathbf{E}\|_F \\ &= \|\tilde{\mathbf{Y}} - \mathbf{X}_{\tilde{\gamma}} \mathbf{X}_{\tilde{\gamma}}^T \tilde{\mathbf{Y}}\|_F + \|\mathbf{E}\|_F \\ &\leq \sqrt{\gamma} \cdot \min_{\mathbf{X} \in \mathcal{X}} \|\tilde{\mathbf{Y}} - \mathbf{X}\mathbf{X}^T \tilde{\mathbf{Y}}\|_F + \|\mathbf{E}\|_F \\ &\leq \sqrt{\gamma} \cdot \|\tilde{\mathbf{Y}} - \mathbf{X}_{\text{opt}} \mathbf{X}_{\text{opt}}^T \tilde{\mathbf{Y}}\|_F + \|\mathbf{E}\|_F \\ &= \sqrt{\gamma} \cdot \|\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T - \mathbf{X}_{\text{opt}} \mathbf{X}_{\text{opt}}^T \tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^T\|_F + \|\mathbf{E}\|_F \\ &= \sqrt{\gamma} \cdot \|\mathbf{Y}\mathbf{Y}^T - \mathbf{E} - \mathbf{X}_{\text{opt}} \mathbf{X}_{\text{opt}}^T (\mathbf{Y}\mathbf{Y}^T - \mathbf{E})\|_F + \|\mathbf{E}\|_F \\ &\leq \sqrt{\gamma} \cdot \|\mathbf{Y}\mathbf{Y}^T - \mathbf{X}_{\text{opt}} \mathbf{X}_{\text{opt}}^T \mathbf{Y}\mathbf{Y}^T\|_F + 2 \cdot \sqrt{\gamma} \cdot \|\mathbf{E}\|_F \\ &= \sqrt{\gamma} \cdot \|\mathbf{Y} - \mathbf{X}_{\text{opt}} \mathbf{X}_{\text{opt}}^T \mathbf{Y}\|_F + 2 \cdot \sqrt{\gamma} \cdot \|\mathbf{E}\|_F \\ &= \sqrt{\gamma} \cdot \sqrt{\mathbf{F}_{\text{opt}}} + 2 \cdot \sqrt{\gamma} \cdot \|\mathbf{E}\|_F \\ &\leq \sqrt{\gamma} \cdot \sqrt{\mathbf{F}_{\text{opt}}} + 2 \cdot \sqrt{\gamma} \cdot \varepsilon \end{aligned}$$

In the above, we used the triangle inequality for the Frobenius norm, the fact that $(\mathbf{I}_n - \mathbf{X}_{\tilde{\gamma}}\mathbf{X}_{\tilde{\gamma}}^T)$ and $(\mathbf{I}_n - \mathbf{X}_{opt}\mathbf{X}_{opt}^T)$ are projection matrices⁵ (combined with the fact that for any projection matrix \mathbf{P} and any matrix \mathbf{Z} : $\|\mathbf{PZ}\|_F \leq \|\mathbf{Z}\|_F$), the fact that for any matrix \mathbf{Q} with orthonormal columns and any matrix \mathbf{X} : $\|\mathbf{XQ}^T\|_F = \|\mathbf{X}\|_F$, the fact that $1 \leq \sqrt{\gamma}$ and the definition of a γ -approximation k -means algorithm.

Overall, we have proved:

$$\|\mathbf{Y} - \mathbf{X}_{\tilde{\gamma}}\mathbf{X}_{\tilde{\gamma}}^T\mathbf{Y}\|_F \leq \sqrt{\gamma} \cdot \sqrt{F_{opt}} + 2 \cdot \sqrt{\gamma} \cdot \varepsilon.$$

Taking squares on both sides in the previous inequality gives:

$$\begin{aligned} \|\mathbf{Y} - \mathbf{X}_{\tilde{\gamma}}\mathbf{X}_{\tilde{\gamma}}^T\mathbf{Y}\|_F^2 &\leq \\ &\leq \gamma \cdot F_{opt} + 4 \cdot \varepsilon \cdot \gamma \sqrt{F_{opt}} + 4 \cdot \gamma \cdot \varepsilon^2 \\ &= \gamma \cdot (F_{opt} + 4 \cdot \varepsilon \sqrt{F_{opt}} + 4 \cdot \varepsilon^2) \end{aligned}$$

Finally, using $\sqrt{F_{opt}} \leq F_{opt}$ shows the claim in the theorem. This bound fails with probability at most $e^{-2n} + 2.35\delta + \delta_\gamma$, which simply follows by taking the union bound on the failure probabilities of Lemma 5 and the γ -approximation k -means algorithm.

Connection to the normalized Laplacian. Towards this end, we need to use a relation between the eigenvalues of $\tilde{\mathbf{W}}$ and the eigenvalues of $\tilde{\mathbf{L}} = \mathbf{D}^{-\frac{1}{2}}\mathbf{L}\mathbf{D}^{-\frac{1}{2}}$. Recall that the Laplacian matrix of the graph is $\mathbf{L} = \mathbf{D} - \mathbf{W}$ and the normalized Laplacian matrix is $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{W}}$. From the last relation, it is easy to see that an eigenvalue of $\tilde{\mathbf{L}}$ equals to 1 minus some eigenvalue of $\tilde{\mathbf{W}}$; the ordering of the eigenvalues, however, is different. Specifically, for $i = 1, 2, \dots, n$, the relation is

$$\lambda_i(\tilde{\mathbf{W}}) = 1 - \lambda_{n-i+1}(\tilde{\mathbf{L}}),$$

where the ordering is:

$$\lambda_1(\tilde{\mathbf{W}}) \geq \lambda_2(\tilde{\mathbf{W}}) \geq \dots \geq \lambda_n(\tilde{\mathbf{W}}),$$

and

$$\lambda_1(\tilde{\mathbf{L}}) \geq \lambda_2(\tilde{\mathbf{L}}) \geq \dots \geq \lambda_n(\tilde{\mathbf{L}}).$$

From this relation:

$$\gamma_k = \frac{\sigma_k(\tilde{\mathbf{W}})}{\sigma_{k+1}(\tilde{\mathbf{W}})} = \frac{1 - \lambda_{n-k+1}(\tilde{\mathbf{L}})}{1 - \lambda_{n-k}(\tilde{\mathbf{L}})}.$$

⁵A matrix \mathbf{P} is called a projection matrix if is square and $\mathbf{P}^2 = \mathbf{P}$.

5. Experiments

To conduct our experiments, we developed high-quality MATLAB versions of the spectral clustering algorithms. In the remainder of this section, we refer to the clustering algorithm in (Shi & Malik, 2000a) as “exact algorithm”. We refer to the modified version that uses the power method as “approximate algorithm”. To measure clustering quality, we used normalized mutual information (Manning et al., 2008):

$$NMI(\Omega; C) = \frac{I(\Omega; C)}{\frac{1}{2}(H(\Omega) + H(C))}.$$

Here,

$$\Omega = \{\omega_1, \omega_2, \dots, \omega_k\}$$

is the set of discovered clusters and

$$C = \{c_1, c_2, \dots, c_k\}$$

is the set of true class labels. Also,

$$I(\Omega; C) = \sum_k \sum_j P(\omega_k \cap c_j) \log \frac{P(\omega_k \cap c_j)}{P(\omega_k)P(c_j)}$$

is the mutual information between Ω and C ,

$$H(\Omega) = - \sum_k P(\omega_k) \log P(\omega_k)$$

and

$$H(C) = - \sum_k P(c_k) \log P(c_k)$$

are the entropies of Ω and C , where P is the probability that is estimated using maximum-likelihood. NMI is a value in $[0, 1]$, where values closer to 1 represent better clusterings. The exact algorithm uses MATLAB’s `svds` function to compute the top- k singular vectors. The approximate algorithm exploits the tall-thin structure of \mathbf{B} and computes $\tilde{\mathbf{Y}}$ using MATLAB’s `svd` function⁶. The approximate algorithm uses MATLAB’s `normrnd` function to generate the random Gaussian matrix \mathbf{S} . We used MATLAB’s `kmeans` function with the options ‘EmptyAction’, ‘singleton’, ‘MaxIter’, 100, ‘Replicates’, 10. All our experiments were run using MATLAB 8.1.0.604 (R2013a) on a 1.4 GHz Intel Core i5 dual-core processor running OS X 10.9.5 with 8GB 1600 MHz DDR3 RAM. Finally, all reported running times are for computing \mathbf{Y} for the exact algorithm and $\tilde{\mathbf{Y}}$ for the approximate algorithm (given $\tilde{\mathbf{W}}$).

5.1. Spectral Clustering Accuracy

We ran our experiments on four multi-class datasets from the libSVMTools webpage (Table 1). To compute \mathbf{W} ,

⁶`[U, S, V] = svd(B' * B); tildeY = B * V * S^(-.5);`

Name	Exact		Approximate				
	NMI	time (secs)	p=2		Best under exact time		
			NMI	time (secs)	NMI	time (secs)	p
SatImage	0.5905	1.270	0.5713	0.310	0.6007	0.690	6
Segment	0.7007	1.185	0.2240	0.126	0.5305	0.530	10
Vehicle	0.1655	0.024	0.2191	0.009	0.2449	0.022	6
Vowel	0.4304	0.016	0.3829	0.003	0.4307	0.005	3

Table 2. Spectral Clustering results for exact and approximate algorithms on the datasets from Table 1. For approximate algorithms, we report two sets of numbers: (1) the NMI achieved at $p = 2$ along with the time and (2) the best NMI achieved while taking less time than the exact algorithm. We see that the approximate algorithm performs *at least* as good as the exact algorithm for all but the Segment dataset.

Name	n	d	#nnz	#classes
SatImage	4435	36	158048	6
Segment	2310	19	41469	7
Vehicle	846	18	14927	4
Vowel	528	10	5280	11

Table 1. The libSVM multi-class classification datasets (Chang & Lin, 2011) used for our spectral clustering experiments.

we use the heat kernel: $\mathbf{W}_{ij} = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{\sigma_{ij}}}$, where $\mathbf{x}_i \in \mathbb{R}^d$ and $\mathbf{x}_j \in \mathbb{R}^d$ are the data points and σ_{ij} is a tuning parameter; σ_{ij} is determined using the self-tuning method described in (Zelnik-Manor & Perona, 2004). That is, for each data point i , \mathbf{x}_i is computed to be the Euclidean distance of the ℓ^{th} furthest neighbor from i ; then σ_{ij} is set to be $\mathbf{x}_i \cdot \mathbf{x}_j$ for every (i, j) ; in our experiments, we report the results for $\ell = 7$. To determine the quality of the clustering obtained by the approximate algorithm and to determine the effect of p (number of power iterations) on the clustering quality and running time, we varied p from 0 to 10. In columns 2 and 3, we see the results for the exact algorithm, which serve as our baseline for quality and performance of the approximate algorithm. In columns 4 and 5, we see the NMI with 2 power iterations ($p = 2$) along with the running time. Immediately, we see that even at $p = 2$, the Vehicle dataset achieves better accuracy than the exact algorithm while resulting in a 2.5x speedup. The only outlier is the Segment dataset, which achieves poor NMI at $p = 2$. In columns 6–8, we report the best NMI that was achieved by the approximate algorithm while staying under the time taken by the exact algorithm; we also report the value of p and the time taken (in seconds) for this result. We see that even when constrained to run in less time than the exact algorithm, the approximate algorithm bests the NMI achieved by the exact algorithm. For example, at $p = 6$, the approximate algorithm reports NMI=0.2449 for Vehicle, as opposed to NMI=0.1655 achieved by the exact algorithm. We can see that in many cases, 2 subspace iterations suffice to get good quality clustering (Segment dataset is the exception).

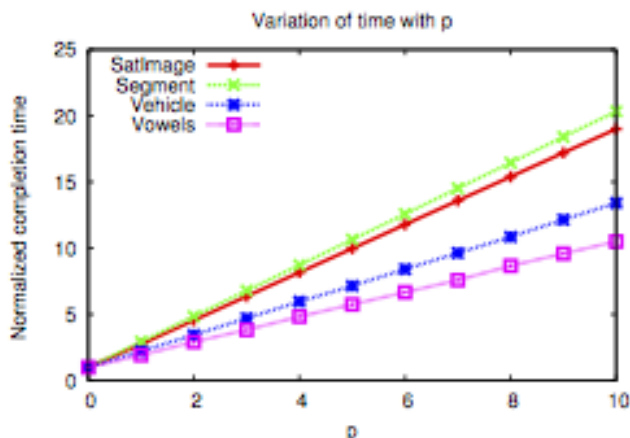


Figure 3. Increase in running time (to compute $\tilde{\mathbf{Y}}$) normalized by running time when $p = 0$. The baseline times (at $p=0$) are SatImage (0.0648 secs), Segment (0.0261 secs), Vehicle (0.0027 secs), and Vowel (0.0012 secs).

Figure 3 depicts the relation between p and the running time (to compute $\tilde{\mathbf{Y}}$) of the approximate algorithm. All the times are normalized by the time taken when $p = 0$ to enable reporting numbers for all datasets on the same scale. As expected, as p increases, we see a linear increase.

Acknowledgment

We would like to thank Yiannis Koutis and James Lee for useful discussions.

References

- Baker, C.T.H. *The numerical treatment of integral equations*, volume 13. Clarendon Press, 1977.
- Bandeira, Afonso S, Singer, Amit, and Spielman, Daniel A. A cheeger inequality for the graph connection laplacian. *SIAM Journal on Matrix Analysis and Applications*, 34(4):1611–1630, 2013.

- Belkin, Mikhail and Niyogi, Partha. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *NIPS*, 2001.
- Bezdek, J.C., Hathaway, R.J., Huband, J.M., Leckie, C., and Kotagiri, R. Approximate clustering in very large relational data. *International journal of intelligent systems*, 21(8):817–841, 2006.
- Boutsidis, Christos and Magdon-Ismail, Malik. Deterministic feature selection for k-means clustering. *Information Theory, IEEE Transactions on*, 59(9):6099–6110, 2013.
- Boutsidis, Christos and Magdon-Ismail, Malik. Faster svd-truncated regularized least-squares. *ISIT*, 2014.
- Chang, Chih-Chung and Lin, Chih-Jen. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Cohen, Michael, E, Sam, Musco, Cameron, Musco, Christopher, and Persu, Madalina. Dimensionality reduction for k-means clustering and low rank approximation. *arXiv preprint arXiv:1410.6801*, 2014.
- Fiedler, Miroslav. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.
- Fowlkes, C., Belongie, S., Chung, F., and Malik, J. Spectral grouping using the Nystrom method. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(2): 214–225, 2004.
- Golub, Gene H and Van Loan, Charles F. *Matrix computations*, volume 3. JHU Press, 2012.
- Halko, Nathan, Martinsson, Per-Gunnar, and Tropp, Joel A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- Kumar, Amit, Sabharwal, Yogish, and Sen, Sandeep. A simple linear time $(1 + \epsilon)$ -approximation algorithm for geometric k-means clustering in any dimensions. In *FOCS*, 2004.
- Lee, James R, Oveis Gharan, Shayan, and Trevisan, Luca. Multi-way spectral partitioning and higher-order cheeger inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 1117–1130. ACM, 2012.
- Lin, Frank and Cohen, William W. Power iteration clustering. In *ICML*, 2010.
- Liu, Rong and Zhang, Hao. Segmentation of 3d meshes through spectral clustering. In *Computer Graphics and Applications, 2004. PG 2004. Proceedings. 12th Pacific Conference on*, pp. 298–305. IEEE, 2004.
- Manning, Christopher D., Raghavan, Prabhakar, and Schütze, Hinrich. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- Ng, Andrew Y, Jordan, Michael I, Weiss, Yair, et al. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 2:849–856, 2002.
- Nyström, E. Über die praktische Auflösung von Integralgleichungen mit Anwendungen auf Randwertaufgaben. *Acta Mathematica*, 54(1):185–204, 1930.
- Ostrovsky, Rafail, Rabani, Yuval, Schulman, Leonard J, and Swamy, Chaitanya. The effectiveness of lloyd-type methods for the k-means problem. In *FOCS*, 2006.
- Pavan, M. and Pelillo, M. Efficient out-of-sample extension of dominant-set clusters. *NIPS*, 2005.
- Shi, J. and Malik, J. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8):888–905, 2000a.
- Shi, Jianbo and Malik, Jitendra. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8):888–905, 2000b.
- Smyth, S and White, S. A spectral clustering approach to finding communities in graphs. In *SDM*, 2005.
- Spielman, Daniel and Teng, Shang-Hua. Nearly-linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *Arxiv preprint, http://arxiv.org/pdf/cs/0607105v4.pdf*, 2009.
- Von Luxburg, Ulrike. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- Wang, L., Leckie, C., Ramamohanarao, K., and Bezdek, J. Approximate spectral clustering. *Advances in Knowledge Discovery and Data Mining*, pp. 134–146, 2009.
- Weiss, Yair. Segmentation using eigenvectors: a unifying view. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pp. 975–982. IEEE, 1999.
- Woodruff, David P. Sketching as a tool for numerical linear algebra. *arXiv preprint arXiv:1411.4357*, 2014.
- Yan, D., Huang, L., and Jordan, M.I. Fast approximate spectral clustering. In *KDD*, 2009.
- Zelnik-Manor, Lihi and Perona, Pietro. Self-tuning spectral clustering. In *NIPS*, 2004.