

Scalable Kernel k -Means Clustering with Nyström Approximation: Relative-Error Bounds

Shusen Wang

*Department of Computer Science
Stevens Institute of Technology
Hoboken, NJ 07030, USA*

SHUSEN.WANG@STEVENS.EDU

Alex Gittens

*Computer Science Department
Rensselaer Polytechnic Institute
Troy, NY 12180, USA*

GITTEA@RPI.EDU

Michael W. Mahoney

*International Computer Science Institute and Department of Statistics
University of California at Berkeley
Berkeley, CA 94720, USA*

MMAHONEY@STAT.BERKELEY.EDU

Editor: Sanjiv Kumar

Abstract

Kernel k -means clustering can correctly identify and extract a far more varied collection of cluster structures than the linear k -means clustering algorithm. However, kernel k -means clustering is computationally expensive when the non-linear feature map is high-dimensional and there are many input points.

Kernel approximation, e.g., the Nyström method, has been applied in previous works to approximately solve kernel learning problems when both of the above conditions are present. This work analyzes the application of this paradigm to kernel k -means clustering, and shows that applying the linear k -means clustering algorithm to $\frac{k}{\epsilon}(1 + o(1))$ features constructed using a so-called rank-restricted Nyström approximation results in cluster assignments that satisfy a $1 + \epsilon$ approximation ratio in terms of the kernel k -means cost function, relative to the guarantee provided by the same algorithm without the use of the Nyström method. As part of the analysis, this work establishes a novel $1 + \epsilon$ relative-error trace norm guarantee for low-rank approximation using the rank-restricted Nyström approximation.

Empirical evaluations on the 8.1 million instance MNIST8M dataset demonstrate the scalability and usefulness of kernel k -means clustering with Nyström approximation. This work argues that spectral clustering using Nyström approximation—a popular and computationally efficient, but theoretically unsound approach to non-linear clustering—should be replaced with the efficient and theoretically sound combination of kernel k -means clustering with Nyström approximation. The superior performance of the latter approach is empirically verified.

Keywords: kernel k -means clustering, the Nyström method, randomized linear algebra

1. Introduction

Cluster analysis divides a data set into several groups using information found only in the data points. Clustering can be used in an exploratory manner to discover meaningful groupings within a data set, or it can serve as the starting point for more advanced analyses. As such, applications of clustering abound in machine learning and data analysis, including, *inter alia*: genetic expression analysis (Sharan et al., 2002), market segmentation (Chaturvedi et al., 1997), social network analysis (Handcock et al., 2007), image segmentation (Haralick and Shapiro, 1985), anomaly detection (Chandola et al., 2009), collaborative filtering (Ungar and Foster, 1998), and fast approximate learning of non-linear models (Si et al., 2014).

Linear k -means clustering is a standard and well-regarded approach to cluster analysis that partitions input vectors $\{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathbb{R}^d$ into k clusters, in an unsupervised manner, by assigning each vector to the cluster with the nearest centroid. Formally, linear k -means clustering seeks to partition the set $[n] = \{1, \dots, n\}$ into k disjoint sets $\mathcal{J}_1, \dots, \mathcal{J}_k$ by solving

$$\operatorname{argmin}_{\mathcal{J}_1, \dots, \mathcal{J}_k} \frac{1}{n} \sum_{i=1}^k \sum_{j \in \mathcal{J}_i} \left\| \mathbf{a}_j - \frac{1}{|\mathcal{J}_i|} \sum_{l \in \mathcal{J}_i} \mathbf{a}_l \right\|_2^2. \quad (1)$$

Lloyd’s algorithm (Lloyd, 1982) is a standard approach for finding local minimizers of (1), and is a staple in data mining and machine learning.

Despite its popularity, linear k -means clustering is not a universal solution to all clustering problems. In particular, linear k -means clustering strongly biases the recovered clusters towards isotropy and sphericity. Applied to the data in Figure 1(a), Lloyd’s algorithm is perfectly capable of partitioning the data into three clusters which fit these assumptions. However, the data in Figure 1(b) do not fit these assumptions: the clusters are ring-shaped and have coincident centers, so minimizing the linear k -means objective does not recover these clusters.




Figure 1: Figures 1(a) and 1(b) show two sets of two-dimensional points, with different colors indicating different clusters. Figure 1(a) is separable using linear k -means; Figure 1(b) is inseparable using linear k -means. Figure 1(c) shows the first two dimensions of the feature vectors $\mathbf{k}_1, \dots, \mathbf{k}_n \in \mathbb{R}^n$ derived from a kernel.

To extend the scope of k -means clustering to include anisotropic, non-spherical clusters such as those depicted in Figure 1(b), Schölkopf et al. (1998) proposed to perform linear k -means clustering in a nonlinear feature space instead of the input space. After choosing a feature function $\phi : \mathbb{R}^d \mapsto \mathcal{F}$ to map the input vectors non-linearly into feature vectors,

they propose minimizing the objective function

$$\operatorname{argmin}_{\mathcal{J}_1, \dots, \mathcal{J}_k} \frac{1}{n} \sum_{i=1}^k \sum_{j \in \mathcal{J}_i} \left\| \phi(\mathbf{a}_j) - \frac{1}{|\mathcal{J}_i|} \sum_{l \in \mathcal{J}_i} \phi(\mathbf{a}_l) \right\|_2^2, \quad (2)$$

where $\{\mathcal{J}_1, \dots, \mathcal{J}_k\}$ denotes a k -partition of $[n]$. The “kernel trick” enables us to minimize this objective without explicitly computing the potentially high-dimensional features, as inner products in feature space can be computed implicitly by evaluating the kernel function

$$\kappa(\mathbf{a}_i, \mathbf{a}_j) = \langle \phi(\mathbf{a}_i), \phi(\mathbf{a}_j) \rangle.$$

Thus the information required to solve the kernel k -means problem (2), is present in the kernel matrix $\mathbf{K} = [\kappa(\mathbf{a}_i, \mathbf{a}_j)]_{ij} \in \mathbb{R}^{n \times n}$.

Let $\mathbf{K} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$ be the full eigenvalue decomposition (EVD) of the kernel matrix and $\mathbf{k}_1, \dots, \mathbf{k}_n \in \mathbb{R}^n$ be the rows of $\mathbf{V}\mathbf{\Lambda}^{1/2} \in \mathbb{R}^{n \times n}$. It can be shown (see Appendix B.3) that the solution of (2) is identical to the solution of

$$\operatorname{argmin}_{\mathcal{J}_1, \dots, \mathcal{J}_k} \frac{1}{n} \sum_{i=1}^k \sum_{j \in \mathcal{J}_i} \left\| \mathbf{k}_j - \frac{1}{|\mathcal{J}_i|} \sum_{l \in \mathcal{J}_i} \mathbf{k}_l \right\|_2^2. \quad (3)$$

To demonstrate the power of kernel k -means clustering, consider the dataset in Figure 1(b). We use the Gaussian RBF kernel

$$\kappa(\mathbf{a}, \mathbf{a}') = \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{a} - \mathbf{a}'\|_2^2\right)$$

with $\sigma = 0.3$, and form the corresponding kernel matrix of the data in Figure 1(b). Figure 1(c) scatterplots the first two coordinates of the feature vectors $\mathbf{k}_1, \dots, \mathbf{k}_n$. Clearly, the first coordinate of the feature vectors already separates the two classes well, so k -means clustering using the non-linear features $\mathbf{k}_1, \dots, \mathbf{k}_n$ has a better chance of separating the two classes.

Although it is more generally applicable than linear k -means clustering, kernel k -means clustering is computationally expensive. As a baseline, we consider the cost of optimizing (3). The formation of the kernel matrix \mathbf{K} given the input vectors $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$ costs $\mathcal{O}(n^2d)$ time. The objective in (3) can then be (approximately) minimized using Lloyd’s algorithm at a cost of $\mathcal{O}(n^2k)$ time per iteration. This requires the n -dimensional non-linear feature vectors obtained from the full EVD of \mathbf{K} ; computing these feature vectors takes $\mathcal{O}(n^3)$ time, because \mathbf{K} is, in general, full-rank. Thus, approximately solving the kernel k -means clustering problem by optimizing (3) costs $\mathcal{O}(n^3 + n^2d + Tn^2k)$ time, where T is the number of iterations of Lloyd’s algorithm.

Kernel approximation techniques, including the Nyström method (Nyström, 1930; Williams and Seeger, 2001; Gittens and Mahoney, 2016) and random feature maps (Rahimi and Recht, 2007), have been applied to decrease the cost of solving kernelized machine learning problems: the idea is to replace \mathbf{K} with a low-rank approximation, which allows for more efficient computations. Chitta et al. (2011, 2012) proposed to apply kernel approximations to efficiently approximate kernel k -means clustering. Although kernel approximations mitigate the computational challenges of kernel k -means clustering, the aforementioned works do not provide guarantees on the clustering performance: how

accurate must the low-rank approximation of \mathbf{K} be to ensure near optimality of the approximate clustering obtained via this method?

We propose a provable approximate solution to the kernel k -means problem based on the Nyström approximation. Our method has three steps: first, extract c ($c \ll n$) features using the Nyström method; second, reduce the features to s dimensions ($k \leq s < c$) using the truncated SVD;¹ third, apply any off-the-shelf linear k -means clustering algorithm upon the s -dimensional features to obtain the final clusters. The total time complexity of the first two steps is $\mathcal{O}(ndc + nc^2)$. The time complexity of the third step depends on the specific linear k -means algorithm; for example, using Lloyd’s algorithm, the per-iteration complexity is $\mathcal{O}(nsk)$, and the number of iterations may depend on s .²

Our method comes with a strong approximation ratio guarantee. Suppose we set $s = k/\epsilon$ and $c = \tilde{\mathcal{O}}(\mu s/\epsilon)$ for any $\epsilon \in (0, 1)$, where $\mu \in [1, \frac{n}{s}]$ is the coherence parameter of the dominant s -dimensional singular space of the kernel matrix \mathbf{K} . Also suppose the standard kernel k -means and our approximate method use the same linear k -means clustering algorithm, e.g., Lloyd’s algorithm or some other algorithm that comes with different provable approximation guarantees. As guaranteed by Theorem 2, when the quality of the clustering is measured by the cost function defined in (2), with probability at least 0.9, our algorithm returns a clustering that is at most $1 + \mathcal{O}(\epsilon)$ times worse than the standard kernel k -means clustering. Our theory makes explicit the trade-off between accuracy and computation: larger s and c lead to high accuracy and also high computational cost.

Spectral clustering (Shi and Malik, 2000; Ng et al., 2002) is a popular alternative to kernel k -means clustering that can also partition non-linearly separable data such as those in Figure 1(b). Unfortunately, because it requires computing an $n \times n$ affinity matrix and the top k eigenvectors of the corresponding graph Laplacian, spectral clustering is inefficient for large n . Fowlkes et al. (2004) applied the Nyström approximation to increase the scalability of spectral clustering. Since then, spectral clustering with Nyström approximation has been used in many works, e.g., (Arikan, 2006; Berg et al., 2004; Chen et al., 2011; Wang et al., 2016b; Weiss et al., 2009; Zhang and Kwok, 2010). Despite its popularity in practice, this approach does not come with guarantees on the approximation ratio for the obtained clustering. Our algorithm, which combines kernel k -means with Nyström approximation, is an equally computationally efficient alternative that comes with strong bounds on the approximation ratio, and can be used wherever spectral clustering is applied.

1.1. Contributions

Using tools developed in (Boutsidis et al., 2015; Cohen et al., 2015; Feldman et al., 2013), we rigorously analyze the performance of approximate kernel k -means clustering with the Nyström approximation, and show that a rank- $\frac{k}{\epsilon}$ Nyström approximation delivers a $1 + \mathcal{O}(\epsilon)$

-
1. This is why our variant is called the rank-restricted Nyström approximation. The rank-restriction serves two purposes. First, although we do not know whether the rank-restriction is necessary for the $1 + \epsilon$ bound, we are unable to establish the bound without it. Second, the rank-restriction makes the third step, linear k -means clustering, much less costly. For the computational benefit, previous works (Boutsidis et al., 2009, 2010, 2015; Cohen et al., 2015; Feldman et al., 2013) have considered dimensionality reduction for linear k -means clustering.
 2. Without the rank-restriction, the per-iteration cost would be $\mathcal{O}(nck)$, and the number of iterations may depend on c .

approximation ratio guarantee, relative to the guarantee provided by the same algorithm without the use of the Nyström method.

As part of the analysis of kernel k -means with Nyström approximation, we establish the first relative-error bound for rank-restricted Nyström approximation,³ which has independent interest.

Kernel k -means clustering and spectral clustering are competing solutions to the nonlinear clustering problem, neither of which scales well with n . Fowlkes et al. (2004) introduced the use of Nyström approximations to make spectral clustering scalable; this approach has become popular in machine learning. We identify fundamental mathematical problems with this heuristic. These concerns and an empirical comparison establish that our proposed combination of kernel k -means with rank-restricted Nyström approximation is a theoretically well-founded and empirically competitive alternative to spectral clustering with Nyström approximation.

Finally, we demonstrate the scalability of this approach by measuring the performance of an Apache Spark implementation of a distributed version of our approximate kernel k -means clustering algorithm using the MNIST8M data set, which has 8.1 million instances and 10 classes.

1.2. Relation to Prior Work

The key to our analysis of the proposed approximate kernel k -means clustering algorithm is a novel relative-error trace norm bound for a *rank-restricted* Nyström approximation. We restrict the rank of the Nyström approximation in a non-standard manner (see Remark 1). Our relative-error trace norm bound is not a simple consequence of the existing bounds for non-rank-restricted Nyström approximation such as the ones provided by Gittens and Mahoney (2016). The relative-error bound which we provide for the rank-restricted Nyström approximation is potentially useful in other applications involving the Nyström method.

The projection-cost preservation (PCP) property (Cohen et al., 2015; Feldman et al., 2013) is an important tool for analyzing approximate linear k -means clustering. We apply our novel relative-error trace norm bound as well as existing tools in (Cohen et al., 2015) to prove that the rank-restricted Nyström approximation enjoys the PCP property. We do not rule out the possibility that the non-rank-restricted (rank- c) Nyström approximation satisfies the PCP property and/or also enjoys a $1 + \epsilon$ approximation ratio guarantee when applied to kernel k -means clustering. However, the cost of the linear k -means clustering step in the algorithm is proportional to the dimensionality of the feature vectors, so the rank-restricted Nyström approximation, which produces s -dimensional feature vectors, where $s < c$, is more computationally desirable.

Musco and Musco (2017) similarly establishes a $1 + \epsilon$ approximation ratio for the kernel k -means objective when a Nyström approximation is used in place of the full kernel matrix. Specifically, Musco and Musco (2017) shows that when $c = \mathcal{O}(\frac{k}{\epsilon} \log \frac{k}{\epsilon})$ columns of \mathbf{K} are sampled using *ridge leverage score (RLS) sampling* (Alaoui and Mahoney, 2015; Cohen et al., 2017; Musco and Musco, 2017) and are used to form a Nyström approximation,

3. Similar relative-error bounds were independently developed by contemporaneous work of Tropp et al. (2017), in service of the analysis of a novel streaming algorithm for fixed-rank approximation of positive semidefinite matrices. Preliminary versions of this work and theirs were simultaneously submitted to arXiv in June 2017.

Table 1: Commonly used parameters. It holds that $k \leq s \leq c \leq n$.

Notation	Definition
n	number of samples
d	number of features (attributes)
k	number of clusters
s	target rank of the Nyström approximation
c	sketch size of the Nyström approximation

then applying linear k -means clustering to the c -dimensional Nyström features returns a clustering that has objective value at most $1 + \epsilon$ times as large as the objective value of the best clustering. Our theory is independent of that in Musco and Musco (2017), and differs in that (1) Musco and Musco (2017) applies specifically to Nyström approximations formed using RLS sampling, whereas our guarantees apply to any sketching method that satisfies the “subspace embedding” and “matrix multiplication” properties (see Lemma 10 for definitions of these two properties); (2) Musco and Musco (2017) establishes a $1 + \epsilon$ approximation ratio for the non-rank-restricted RLS-Nyström approximation, whereas we establish a $1 + \epsilon$ approximation ratio for the (more computationally efficient) rank-restricted Nyström approximation.

1.3. Paper Organization

In Section 2, we start with a definition of the notation used throughout this paper as well as a background on matrix sketching methods. Then, in Section 3, we present our main theoretical results: Section 3.1 presents an improved relative-error rank-restricted Nyström approximation; Section 3.2 presents the main theoretical results on kernel k -means with Nyström approximation; and Section 3.3 studies kernel k -means with kernel principal component analysis. Section 4 discusses and evaluates the theoretical and empirical merits of kernel k -means clustering versus spectral clustering, when each is approximated using Nyström approximation. Section 5 empirically compares the Nyström method and the random feature maps for the kernel k -means clustering on medium-scale data. Section 6 presents a large-scale distributed implementation in Apache Spark and its empirical evaluation on a data set with 8.1 million points. Section 7 provides a brief conclusion. Proofs are provided in the Appendices.

2. Notation

This section defines the notation used throughout this paper. A set of commonly used parameters is summarized in Table 1.

Matrices and vectors. We take \mathbf{I}_n to be the $n \times n$ identity matrix, $\mathbf{0}$ to be a vector or matrix of all zeros of the appropriate size, and $\mathbf{1}_n$ to be the n -dimensional vector of all ones.

Sets. The set $\{1, 2, \dots, n\}$ is written as $[n]$. We call $\{\mathcal{J}_1, \dots, \mathcal{J}_k\}$ a k -partition of $[n]$ if $\mathcal{J}_1 \cup \dots \cup \mathcal{J}_k = [n]$ and $\mathcal{J}_p \cap \mathcal{J}_q = \emptyset$ when $p \neq q$. Let $|\mathcal{J}|$ denote the cardinality of the set \mathcal{J} .

Singular value decomposition (SVD). Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\rho = \text{rank}(\mathbf{A})$. A (compact) singular value decomposition (SVD) is defined by

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_{i=1}^{\rho} \sigma_i(\mathbf{A})\mathbf{u}_i\mathbf{v}_i^T,$$

where \mathbf{U} , $\mathbf{\Sigma}$, \mathbf{V} are an $n \times \rho$ column-orthogonal matrix, a $\rho \times \rho$ diagonal matrix with nonnegative entries, and a $d \times \rho$ column-orthogonal matrix, respectively. If \mathbf{A} is symmetric positive semi-definite (SPSD), then $\mathbf{U} = \mathbf{V}$, and this decomposition is also called the (reduced) eigenvalue decomposition (EVD). By convention, we take $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_{\rho}(\mathbf{A})$.

Truncated SVD. The matrix $\mathbf{A}_s = \sum_{i=1}^s \sigma_i(\mathbf{A})\mathbf{u}_i\mathbf{v}_i^T$ is a rank- s truncated SVD of \mathbf{A} , and is an optimal rank- s approximation to \mathbf{A} when the approximation error is measured in a unitarily invariant norm.

The Moore-Penrose inverse of \mathbf{A} is defined by $\mathbf{A}^\dagger = \mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^T$.

Leverage score and coherence. Let $\mathbf{U} \in \mathbb{R}^{n \times \rho}$ be defined in the above and \mathbf{u}_i be the i -th row of \mathbf{U} . The row leverage scores of \mathbf{A} are $l_i = \|\mathbf{u}_i\|_2^2$ for $i \in [n]$. The row coherence of \mathbf{A} is $\mu(\mathbf{A}) = \frac{n}{\rho} \max_i \|\mathbf{u}_i\|_2^2$. The leverage scores for a matrix \mathbf{A} can be computed exactly in the time it takes to compute the matrix \mathbf{U} ; and the leverage scores can be approximated (in theory (Drineas et al., 2012) and in practice (Gittens and Mahoney, 2016)) in roughly the time it takes to apply a random projection matrix to the matrix \mathbf{A} .

Matrix norms. We use three matrix norms in this paper:

$$\begin{aligned} \text{Frobenius Norm: } \quad \|\mathbf{A}\|_F &= \sqrt{\sum_{i,j} a_{ij}^2} = \sqrt{\sum_i \sigma_i^2(\mathbf{A})}; \\ \text{Spectral Norm: } \quad \|\mathbf{A}\|_2 &= \max_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2 = \sigma_1(\mathbf{A}); \\ \text{Trace Norm: } \quad \|\mathbf{A}\|_* &= \sum_i \sigma_i(\mathbf{A}). \end{aligned}$$

Any square matrix satisfies $\text{tr}(\mathbf{A}) \leq \|\mathbf{A}\|_*$. If additionally \mathbf{A} is SPSP, then $\text{tr}(\mathbf{A}) = \|\mathbf{A}\|_*$.

Matrix sketching Here, we briefly review matrix sketching methods that are commonly used within randomized linear algebra (RLA) (Mahoney, 2011).

Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we call $\mathbf{C} = \mathbf{A}\mathbf{P} \in \mathbb{R}^{m \times c}$ (typically $c \ll n$) a *sketch* of \mathbf{A} and $\mathbf{P} \in \mathbb{R}^{n \times c}$ a *sketching matrix*. Within RLA, sketching has emerged as a powerful primitive, where one is primarily interested in using random projections and random sampling to construct randomized sketches (Mahoney, 2011; Drineas and Mahoney, 2016). In particular, sketching is useful as it allows large matrices to be replaced with smaller matrices which are more amenable to efficient computation, but provably retain almost optimal accuracy in many computations (Mahoney, 2011; Woodruff, 2014). The columns of \mathbf{C} typically comprise a rescaled subset of the columns of \mathbf{A} , or random linear combinations of the columns of \mathbf{A} ; the former type of sketching is called *column selection* or *random sampling*, and the latter is referred to as *random projection*.

Column selection forms $\mathbf{C} \in \mathbb{R}^{m \times c}$ using a randomly sampled and rescaled subset of the columns of $\mathbf{A} \in \mathbb{R}^{m \times n}$. Let $p_1, \dots, p_n \in (0, 1)$ be the sampling probabilities associated

with the columns of \mathbf{A} (so that, in particular, $\sum_{i=1}^n p_i = 1$). The columns of the sketch are selected identically and independently as follows: each column of \mathbf{C} is randomly sampled from the columns of \mathbf{A} according to the sampling probabilities and rescaled by $\frac{1}{\sqrt{cp_i}}$, where i is the index of the column of \mathbf{A} that was selected. In our matrix multiplication formulation for sketching, column selection corresponds to a sketching matrix $\mathbf{P} \in \mathbb{R}^{n \times c}$ that has exactly one non-zero entry in each column, whose position and magnitude correspond to the index of the column selected from \mathbf{A} . *Uniform sampling* is column sampling with $p_1 = \dots = p_n = \frac{1}{n}$, and *leverage score sampling* takes $p_i = \frac{l_i}{\sum_j l_j}$ for $i \in [n]$, where l_i is the i -th leverage score of some matrix (typically \mathbf{A} , \mathbf{A}_k , or an randomized approximation thereto) (Drineas et al., 2012).

Gaussian projection is a type of random projection where the sketching matrix is taken to be $\mathbf{P} = \frac{1}{\sqrt{c}}\mathbf{G} \in \mathbb{R}^{n \times c}$; here the entries of \mathbf{G} are i.i.d. $\mathcal{N}(0, 1)$ random variables. Gaussian projection is inefficient relative to column sampling: the formation of a Gaussian sketch of a dense $m \times n$ matrix requires $\mathcal{O}(mnc)$ time. The *Subsampled Randomized Hadamard Transform (SRHT)* is a more efficient alternative that enjoys similar properties to the Gaussian projection (Drineas et al., 2011; Lu et al., 2013; Tropp, 2011), and can be applied to a dense $m \times n$ matrix in only $\mathcal{O}(mn \log c)$ time. The *CountSketch* is even more efficient: it can be applied to any matrix \mathbf{A} in $\mathcal{O}(\text{nnz}(\mathbf{A}))$ time (Clarkson and Woodruff, 2013; Meng and Mahoney, 2013; Nelson and Nguyễn, 2013), where $\text{nnz}(\cdot)$ denotes the number of nonzero entries in a matrix.

3. Our Main Results: Improved SPSD Matrix Approximation and Kernel k -means Approximation

In this section, we present our main theoretical results. We start, in Section 3.1, by presenting Theorem 1, a novel result on SPSD matrix approximation with the rank-restricted Nyström method. This result is of independent interest, and so we present it in detail, but in this paper we will use it to establish our main result. Then, in Section 3.2, we present Theorem 2, which is our main result for approximate kernel k -means with the Nyström approximation. In Section 3.3, we establish novel guarantees on kernel k -means with dimensionality reduction.

3.1. The Nyström Method

The Nyström method (Nyström, 1930) is the most popular kernel approximation method in the machine learning community. Let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be an SPSD matrix and $\mathbf{P} \in \mathbb{R}^{n \times c}$ be a sketching matrix. The Nyström method approximates \mathbf{K} with $\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T$, where $\mathbf{C} = \mathbf{K}\mathbf{P}$ and $\mathbf{W} = \mathbf{P}^T\mathbf{K}\mathbf{P}$. The Nyström method was introduced to the machine learning community by Williams and Seeger (2001); since then, numerous works have studied its theoretical properties, e.g., (Drineas and Mahoney, 2005; Gittens and Mahoney, 2016; Jin et al., 2013; Kumar et al., 2012; Wang and Zhang, 2013; Yang et al., 2012).

Empirical results in (Gittens and Mahoney, 2016; Wang et al., 2016b; Yang et al., 2012) demonstrated that the accuracy of the Nyström method significantly increases when the spectrum of \mathbf{K} decays fast. This suggests that the Nyström approximation captures the dominant eigenspaces of \mathbf{K} , and that error bounds comparing the accuracy of the Nyström

Table 2: Sufficient sketch sizes for several sketching methods used to construct Nyström approximations to the matrix \mathbf{K} . Let $\mathbf{K}_s = \mathbf{V}_s \boldsymbol{\Sigma}_s \mathbf{V}_s^T$ be the truncated SVD of \mathbf{K} , then μ denotes the row coherence of $\mathbf{V}_s \in \mathbb{R}^{n \times s}$ and the leverage score sampling is done using the row leverage scores of \mathbf{V}_s .

sketching	sketch size (c)	time complexity (T)
uniform sampling	$\mathcal{O}(\mu s / \epsilon + \mu s \log s)$	$\mathcal{O}(nc)$
leverage sampling	$\mathcal{O}(s / \epsilon + s \log s)$	$\tilde{\mathcal{O}}(n^2 s)$
Gaussian projection	$\mathcal{O}(s / \epsilon)$	$\mathcal{O}(n^2 c)$
SRHT	$\mathcal{O}((s + \log n)(\epsilon^{-1} + \log s))$	$\mathcal{O}(n^2 \log c)$
CountSketch	$\mathcal{O}(s / \epsilon + s^2)$	$\mathcal{O}(n^2)$

approximation of \mathbf{K} to that of the best rank- s approximation \mathbf{K}_s (for $s < c$) would provide a meaningful measure of the performance of the Nyström kernel approximation. Gittens and Mahoney (2016) established the first relative-error bounds showing that for sufficiently large c , the trace norm error $\|\mathbf{K} - \mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T\|_*$ is comparable to $\|\mathbf{K} - \mathbf{K}_s\|_*$. Such results quantify the benefits of spectral decay to the performance of the Nyström method, and are sufficient to analyze the performance of Nyström approximations in applications such as kernel ridge regression (Alaoui and Mahoney, 2015; Bach, 2013) and kernel support vector machines (Cortes et al., 2010).

However, Gittens and Mahoney (2016) did not analyze the performance of rank-restricted Nyström approximations; they compared the approximation accuracies of the rank- c matrix $\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T$ and the rank- s matrix \mathbf{K}_s (recall $s < c$). In our application to approximate kernel k -means clustering, it is the rank- s matrix $(\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T)_s$ that is of relevance. Given \mathbf{C} and \mathbf{W} , the truncated SVD $((\mathbf{W}^\dagger)^{1/2} \mathbf{C}^T)_s = \tilde{\mathbf{U}}_s \tilde{\boldsymbol{\Sigma}}_s \tilde{\mathbf{V}}_s^T$ can be found using $\mathcal{O}(nc^2)$ time. Then the rank- s Nyström approximation can be written as

$$(\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T)_s = (\mathbf{C}(\mathbf{W}^\dagger)^{1/2})_s (\mathbf{C}(\mathbf{W}^\dagger)^{1/2})_s^T = \tilde{\mathbf{V}}_s \tilde{\boldsymbol{\Sigma}}_s^2 \tilde{\mathbf{V}}_s^T. \quad (4)$$

Theorem 1 provides a relative-error trace norm approximation guarantee for the sketch (4) and is novel; a proof is provided in Appendix A.

Theorem 1 (Relative-Error Rank-Restricted Nyström Approximation) *Let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be an SPSD matrix, s be the target rank, and $\epsilon \in (0, 1)$ be an error parameter. Let $\mathbf{P} \in \mathbb{R}^{n \times c}$ be a sketching matrix corresponding to one of the sketching methods listed in Table 2. Let $\mathbf{C} = \mathbf{K}\mathbf{P}$ and $\mathbf{W} = \mathbf{P}^T \mathbf{K}\mathbf{P}$. Then*

$$\|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T)_s\|_* \leq (1 + \epsilon) \|\mathbf{K} - \mathbf{K}_s\|_*$$

holds with probability at least 0.9. In addition, there exists an $n \times s$ column orthogonal matrix \mathbf{Q} such that $(\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T)_s = \mathbf{K}^{1/2} \mathbf{Q}\mathbf{Q}^T \mathbf{K}^{1/2}$.

Remark 1 (Rank Restrictions) *The traditional rank-restricted Nyström approximation, $\mathbf{C}(\mathbf{W}_s)^\dagger \mathbf{C}^T$, (Drineas and Mahoney, 2005; Fowlkes et al., 2004; Gittens and Mahoney, 2016; Li et al., 2015) is not known to satisfy a relative-error bound of the form*

guaranteed in Theorem 1. Pourkamali-Anaraki and Becker (2016) pointed out the drawbacks of the traditional rank-restricted Nyström approximation and proposed the use of the rank-restricted Nyström approximation $(\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s$ in applications requiring kernel approximations, but provided only empirical evidence of its performance. This work provides guarantees on the approximation error of the rank-restricted Nyström approximation $(\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s$, and applies this approximation to the kernel k -means clustering problem. The contemporaneous work Tropp et al. (2017) provides similar guarantees on the approximation error of $(\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s$, and uses this Nyström approximation as the basis of a streaming algorithm for fixed-rank approximation of positive-semidefinite matrices.

3.2. Main Result for Approximate Kernel k -means

In this section we establish the approximation ratio guarantees for the objective function of kernel k -means clustering. We first define γ -approximate k -means algorithms (where $\gamma \geq 1$), then present our main result in Theorem 2.

Let \mathbf{A} be a matrix with n rows $\mathbf{a}_1, \dots, \mathbf{a}_n$. The objective function for linear k -means clustering over the rows of \mathbf{A} is

$$f(\mathcal{J}_1, \dots, \mathcal{J}_k; \mathbf{A}) = \frac{1}{n} \sum_{i=1}^k \sum_{j \in \mathcal{J}_i} \left\| \mathbf{a}_j - \frac{1}{|\mathcal{J}_i|} \sum_{l \in \mathcal{J}_i} \mathbf{a}_l \right\|_2^2.$$

The minimization of f w.r.t. the k -partition $\{\mathcal{J}_1, \dots, \mathcal{J}_k\}$ is NP-hard (Garey et al., 1982; Aloise et al., 2009; Dasgupta and Freund, 2009; Mahajan et al., 2009; Awasthi et al., 2015), but approximate solutions can be obtained in polynomial time. γ -approximate algorithms capture one useful notion of approximation.

Definition 1 (γ -Approximate Algorithms) *A linear k -means clustering algorithm \mathcal{A}_γ takes as input a matrix \mathbf{Z} with n rows and outputs $\{\mathcal{J}'_1, \dots, \mathcal{J}'_k\}$. We call \mathcal{A}_γ a γ -approximate algorithm if, for any such matrix \mathbf{Z} ,*

$$f(\mathcal{J}'_1, \dots, \mathcal{J}'_k; \mathbf{Z}) \leq \gamma \cdot \min_{\mathcal{J}_1, \dots, \mathcal{J}_k} f(\mathcal{J}_1, \dots, \mathcal{J}_k; \mathbf{Z}).$$

Here $\{\mathcal{J}_1, \dots, \mathcal{J}_k\}$ and $\{\mathcal{J}'_1, \dots, \mathcal{J}'_k\}$ are k -partitions of $[n]$.

Many $(1 + \epsilon)$ -approximation algorithms have been proposed, but they are computationally expensive (Chen, 2009; Har-Peled and Mazumdar, 2004; Kumar et al., 2004; Matousek, 2000). There are also relatively efficient constant factor approximate algorithms, e.g., (Arthur and Vassilvitskii, 2007; Kanungo et al., 2002; Song and Rajasekaran, 2010).

Let ϕ be a feature map, Φ be the matrix with rows $\phi(\mathbf{a}_1), \dots, \phi(\mathbf{a}_n)$, and $\mathbf{K} = \Phi\Phi^T \in \mathbf{R}^{n \times n}$ be the associated kernel matrix. Analogously, we denote the objective function for kernel k -means clustering by

$$f(\mathcal{J}_1, \dots, \mathcal{J}_k; \Phi) = \frac{1}{n} \sum_{i=1}^k \sum_{j \in \mathcal{J}_i} \left\| \phi(\mathbf{a}_j) - \frac{1}{|\mathcal{J}_i|} \sum_{l \in \mathcal{J}_i} \phi(\mathbf{a}_l) \right\|_2^2,$$

where $\{\mathcal{J}_1, \dots, \mathcal{J}_k\}$ is a k -partition of $[n]$.

Theorem 2 (Kernel k -Means with Nyström Approximation) Choose a sketching matrix $\mathbf{P} \in \mathbb{R}^{n \times c}$ and sketch size c consistent with Table 2. Let $\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T$ be the previously defined Nyström approximation of \mathbf{K} . Let $\mathbf{B} \in \mathbb{R}^{n \times s}$ be any matrix satisfying $\mathbf{B}\mathbf{B}^T = (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s$. Let the k -partition $\{\tilde{\mathcal{J}}_1, \dots, \tilde{\mathcal{J}}_k\}$ be the output of a γ -approximate algorithm applied to the rows of \mathbf{B} . With probability at least 0.9,

$$f(\tilde{\mathcal{J}}_1, \dots, \tilde{\mathcal{J}}_k; \Phi) \leq \gamma(1 + \epsilon + \frac{k}{s}) \cdot \min_{\mathcal{J}_1, \dots, \mathcal{J}_k} f(\mathcal{J}_1, \dots, \mathcal{J}_k; \Phi).$$

Remark 2 Kernel k -means clustering is an NP-hard problem. Therefore, instead of comparing with $\min f$, we compare with clusterings obtained using γ -approximate algorithms. Theorem 2 shows that, when uniform sampling to form the Nyström approximation, if $s = \mathcal{O}(\frac{k}{\epsilon})$ and $c = \tilde{\mathcal{O}}(\frac{\mu s}{\epsilon})$, then the returned clustering has an objective value that is at most a factor of ϵ larger than the objective value of the kernel k -means clustering returned by the γ -approximate algorithm.

Remark 3 Assume we are in a practical setting where c , the budget of column samples one can use to form a Nyström approximation, and k , the number of desired cluster centers, are fixed. The pertinent question is how to choose s to produce a high-quality approximate clustering. Theorem 2 shows that for uniform sampling, the error ratio is

$$1 + \epsilon + \frac{k}{s} = 1 + \tilde{\mathcal{O}}(\frac{s\mu}{c}) + \frac{k}{s}.$$

To balance the two sources of error, s must be larger than k , but not too large a fraction of c . To minimize the above error ratio, s should be selected on the order of $\sqrt{kc/\mu}$. Since the matrix coherence μ (≥ 1) is unknown, it can be heuristically treated as a constant.

We empirically study the effect of the values of c and s using a data set comprising 8.1 million samples. Note that computing the kernel k -means clustering objective function requires the formation of the entire kernel matrix \mathbf{K} , which is infeasible for a data set of this size; instead, we use normalized mutual information (NMI) (Strehl and Ghosh, 2002)—a standard measure of the performance of clustering algorithms—to measure the quality of the clustering obtained by approximating kernel k -means clustering using Nyström approximations formed through uniform sampling. NMI scores range from zero to one, with a larger score indicating better performance. We report the results in Figure 2. The complete details of the experiments, including the experimental setting and time costs, are given in Section 6.

From Figure 2(a) we observe that larger values of c lead to better and more stable clusterings: the mean of the NMI increases and its standard deviation decreases. This is reasonable and in accordance with our theory. However, larger values of c incur more computations, so one should choose c to trade off computation and accuracy.

Figure 2(b) shows that for fixed k and c , the clustering performance is not monotonic in s , which matches Theorem 2 (see the discussion in Remark 3). Setting s as small as k results in poor performance. Setting s over-large not only incurs more computations, but also negatively affects clustering performance; this may suggest the necessity of rank-restriction. Furthermore, in this example, $\sqrt{kc} = 126.5$, which corroborates the suggestion made in Remark 3 that setting s around $\sqrt{kc/\mu}$ (where μ is unknown but can be treated as a constant larger than 1) can be a good choice.




Figure 2: Performance of approximate kernel k -means clustering on the MNIST8M data set, for which $n = 8.1 \times 10^6$, $d = 784$, and $k = 10$. We use Nyström approximations formed using uniform sampling, and the built-in k -means algorithm of Spark MLlib (a parallelized variant of the k -means++ method) with the setting `MaxIter = 100`.

Remark 4 Musco and Musco (2017) established a $1 + \epsilon$ approximation ratio for the kernel k -means objective value when a non-rank-restricted Nyström approximation is formed using ridge leverage scores (RLS) sampling; their analysis is specific to RLS sampling and does not extend to other sketching methods. By way of comparison, our analysis covers several popular sampling schemes and applies to rank-restricted Nyström approximations, but does not extend to RLS sampling.

3.3. Approximate Kernel k -Means with KPCA and Power Method

The use of dimensionality reduction to increase the computational efficiency of k -means clustering has been widely studied, e.g. in (Boutsidis et al., 2010, 2015; Cohen et al., 2015; Feldman et al., 2013; Zha et al., 2002). Kernel principal component analysis (KPCA) is particularly well-suited to this application (Dhillon et al., 2004; Ding et al., 2005). Applying Lloyd’s algorithm on the rows of Φ or $\mathbf{K}^{-1/2}$ has an $\mathcal{O}(n^2k)$ per-iteration complexity; if s features are extracted using KPCA and Lloyd’s algorithm is applied to the resulting s -dimensional feature map, then the per-iteration cost reduces to $\mathcal{O}(nsk)$. Proposition 3 states that, to obtain a $1 + \epsilon$ approximation ratio in terms of the kernel k -means objective function, it suffices to use $s = \frac{k}{\epsilon}$ KPCA features. This proposition is a simple consequence of (Cohen et al., 2015).

Proposition 3 (KPCA) Let Φ be a matrix with n rows, and $\mathbf{K} = \Phi\Phi^T \in \mathbb{R}^{n \times n}$ be the corresponding kernel matrix. Let $\mathbf{K}_s = \mathbf{V}_s\mathbf{\Lambda}_s\mathbf{V}_s^T$ be the truncated SVD of \mathbf{K} and take $\mathbf{B} = \mathbf{V}_s\mathbf{\Lambda}_s^{1/2} \in \mathbb{R}^{n \times s}$. Let the k -partition $\{\tilde{\mathcal{J}}_1, \dots, \tilde{\mathcal{J}}_k\}$ be the output of a γ -approximate algorithm applied to the rows of \mathbf{B} . Then

$$f(\tilde{\mathcal{J}}_1, \dots, \tilde{\mathcal{J}}_k; \Phi) \leq \gamma \left(1 + \frac{k}{s}\right) \cdot \min_{\mathcal{J}_1, \dots, \mathcal{J}_k} f(\mathcal{J}_1, \dots, \mathcal{J}_k; \Phi).$$

Algorithm 1 Approximate Kernel k -Means using the Power Method.

- 1: **Input:** kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$, number of clusters k , target dimension $s (\geq k)$, sketch size $c (\geq s)$, number of iterations $t (\geq 1)$
 - 2: Draw a Gaussian projection matrix $\mathbf{P} \in \mathbb{R}^{n \times c}$;
 - 3: **for** all $j \in [t]$ **do**
 - 4: $\mathbf{P} \leftarrow \mathbf{K}\mathbf{P}$;
 - 5: **end for**
 - 6: Orthogonalize \mathbf{P} to obtain $\mathbf{U} \in \mathbb{R}^{n \times c}$;
 - 7: Compute $\mathbf{C} = \mathbf{K}\mathbf{U}$ and $\mathbf{W} = \mathbf{U}^T \mathbf{K}\mathbf{U}$;
 - 8: Compute a $\mathbf{B} \in \mathbb{R}^{n \times s}$ satisfying $(\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T)_s = \mathbf{B}\mathbf{B}^T$;
 - 9: **return** \mathbf{B} .
-

In practice, the truncated SVD (equivalently EVD) of \mathbf{K} is computed using the power method or Krylov subspace methods. These numerical methods do not compute the exact decomposition $\mathbf{K}_s = \mathbf{V}_s \mathbf{\Lambda}_s \mathbf{V}_s^T$, so Proposition 3 is not directly applicable. It is useful to have a theory that captures the effect of realistically inaccurate estimates like $\mathbf{K}_s \approx \tilde{\mathbf{V}}_s \tilde{\mathbf{\Lambda}}_s \tilde{\mathbf{V}}_s^T$ on the clustering process. As one particular example, consider that general-purpose implementations of the truncated SVD attempt to mitigate the fact that the computed decompositions are inaccurate by returning very high-precision solutions, e.g. solutions that satisfy $\|(\mathbf{I}_n - \mathbf{V}_s \mathbf{V}_s^T) \tilde{\mathbf{V}}_s\|_2 \leq 10^{-10}$. Understanding the trade-off between the precision of the truncated SVD solution and the impact on the approximation ratio of the approximate kernel k -means solution allows us to more precisely manage the computational complexity of our algorithms. Are such high-precision solutions necessary for kernel k -means clustering?

Theorem 4 answers this question by establishing that highly accurate eigenspaces are not significantly more useful in approximate kernel k -means clustering than eigenspace estimates with lower accuracy. A low-precision solution obtained by running the power iteration for a few rounds suffices for kernel k -means clustering applications. We prove Theorem 4 in Appendix C.

Theorem 4 (The Power Method) *Let Φ be a matrix with n rows, $\mathbf{K} = \Phi \Phi^T \in \mathbb{R}^{n \times n}$ be the corresponding kernel matrix, and σ_i be the i -th singular value of \mathbf{K} . Fix an error parameter $\epsilon \in (0, 1)$. Run Algorithm 1 with $t = \mathcal{O}(\frac{\log(n/\epsilon)}{\log(\sigma_s/\sigma_{s+1})})$ to obtain $\mathbf{B} \in \mathbb{R}^{n \times s}$. Let the k -partition $\{\tilde{\mathcal{J}}_1, \dots, \tilde{\mathcal{J}}_k\}$ be the output of a γ -approximate algorithm applied to the rows of \mathbf{B} . If $c = s + \mathcal{O}(\log \frac{1}{\delta})$, then*

$$f(\tilde{\mathcal{J}}_1, \dots, \tilde{\mathcal{J}}_k; \Phi) \leq \gamma \left(1 + \epsilon + \frac{k}{s}\right) \cdot \min_{\mathcal{J}_1, \dots, \mathcal{J}_k} f(\mathcal{J}_1, \dots, \mathcal{J}_k; \Phi).$$

holds with probability at least $1 - \delta$. If $c = s$, then the above inequality holds with probability $0.9 - \mathcal{O}(s^{-\tau})$, where τ is a positive constant (Tao and Vu, 2010).

Note that the power method requires forming the entire kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$, which may not fit in memory even in a distributed setting. Therefore, in practice, the power method may not be as efficient as the Nyström approximation with uniform sampling, which avoids forming \mathbf{K} .

Theorem 2, Proposition 3, and Theorem 4 are highly interesting from a theoretical perspective. These results demonstrate that $s = \frac{k}{\epsilon}(1 + o(1))$ features are sufficient to

ensure a $(1 + \epsilon)$ approximation ratio. Prior work (Dhillon et al., 2004; Ding et al., 2005) set $s = k$ and did not provide approximation ratio guarantees. Indeed, a lower bound in the linear k -means clustering case due to (Cohen et al., 2015) shows that $s = \Omega(\frac{k}{\epsilon})$ is necessary to obtain a $1 + \epsilon$ approximation ratio.

4. Comparison to Spectral Clustering with Nyström Approximation

In this section, we provide a brief discussion and empirical comparison of our clustering algorithm, which uses the Nyström method to approximate kernel k -means clustering, with the popular alternative algorithm that uses the Nyström method to approximate spectral clustering.

4.1. Background

Spectral clustering is a method with a long history (Cheeger, 1969; Donath and Hoffman, 1972, 1973; Fiedler, 1973; Guattery and Miller, 1995; Spielman and Teng, 1996). Within machine learning, spectral clustering is more widely used than kernel k -means clustering (Ng et al., 2002; Shi and Malik, 2000), and the use of the Nyström method to speed up spectral clustering has been popular since Fowlkes et al. (2004). Both spectral clustering and kernel k -means clustering can be approximated in time linear in n by using the Nyström method with uniform sampling. Practitioners reading this paper may ask:

How does the approximate kernel k -means clustering algorithm presented here, which uses Nyström approximation, compare to the popular heuristic of combining spectral clustering with Nyström approximation?

Based on our theoretical results and empirical observations, our answer to this reader is:

Although they have equivalent computational costs, kernel k -means clustering with Nyström approximation is both more theoretically sound and more effective in practice than spectral clustering with Nyström approximation.

We first formally describe spectral clustering, and then substantiate our claim regarding the theoretical advantage of our approximate kernel k -means method. Our discussion is limited to the normalized and symmetric graph Laplacians used in Fowlkes et al. (2004), but spectral clustering using asymmetric graph Laplacians encounters similar issues.

4.2. Spectral Clustering with Nyström Approximation

The input to the spectral clustering algorithm is an affinity matrix $\mathbf{K} \in \mathbb{R}_+^{n \times n}$ that measures the pairwise similarities between the points being clustered; typically \mathbf{K} is a kernel matrix or the adjacency matrix of a weighted graph constructed using the data points as vertices. Let $\mathbf{D} = \text{diag}(\mathbf{K}\mathbf{1}_n)$ be the diagonal degree matrix associated with \mathbf{K} , and $\mathbf{L} = \mathbf{I}_n - \mathbf{D}^{-1/2}\mathbf{K}\mathbf{D}^{-1/2}$ be the associated normalized graph Laplacian matrix. Let $\mathbf{V}_k \in \mathbb{R}^{n \times k}$ denote the bottom k eigenvectors of \mathbf{L} , or equivalently, the top k eigenvectors of $\mathbf{D}^{-1/2}\mathbf{K}\mathbf{D}^{-1/2}$. Spectral clustering groups the data points by performing linear k -means clustering on the normalized rows of \mathbf{V}_k . Fowlkes et al. (2004) popularized the application of the Nyström approximation to spectral clustering. This algorithm computes an approximate spectral

clustering by: (1) forming a Nyström approximation to \mathbf{K} , denoted by $\tilde{\mathbf{K}}$; (2) computing the degree matrix $\tilde{\mathbf{D}} = \text{diag}(\tilde{\mathbf{K}}\mathbf{1}_n)$ of $\tilde{\mathbf{K}}$; (3) computing the top k singular vectors $\tilde{\mathbf{V}}_k$ of $\tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{K}}\tilde{\mathbf{D}}^{-1/2}$, which are equivalent to the bottom k eigenvectors of $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{K}}\tilde{\mathbf{D}}^{-1/2}$; (4) performing linear k -means over the normalized rows of $\tilde{\mathbf{V}}_k \in \mathbb{R}^{n \times k}$.

To the best of our knowledge, spectral clustering with Nyström approximation does not have a bounded approximation ratio relative to exact spectral clustering. In fact, it seems unlikely that the approximation ratio could be bounded, as there are fundamental problems with the application of the Nyström approximation to the affinity matrix.

- The affinity matrix \mathbf{K} used in spectral clustering must be elementwise nonnegative. However, the Nyström approximation of such a matrix can have numerous negative entries, so $\tilde{\mathbf{K}}$ is, in general, not proper input for the spectral clustering algorithm. In particular, the approximated degree matrix $\tilde{\mathbf{D}} = \text{diag}(\tilde{\mathbf{K}}\mathbf{1}_n)$ may have negative diagonal entries, so $\tilde{\mathbf{D}}^{-1/2}$ is not guaranteed to be a real matrix; such exceptions must be handled heuristically. The approximate asymmetric Laplacian $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{K}}\tilde{\mathbf{D}}^{-1}$ does avoid the introduction of complex values; however, the negative entries in $\tilde{\mathbf{D}}^{-1}$ negate whole columns of $\tilde{\mathbf{K}}$, leading to less meaningful negative similarities/distances.
- Even if $\tilde{\mathbf{D}}^{-1/2}$ is real, the matrix $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{K}}\tilde{\mathbf{D}}^{-1/2}$ may not be SPSD, much less a Laplacian matrix. Thus the bottom eigenvectors of $\tilde{\mathbf{L}}$ cannot be viewed as useful coordinates for linear k -means clustering in the same way that the eigenvectors of \mathbf{L} can be.
- Such approximation is also problematic in terms of matrix approximation accuracy. Even when $\tilde{\mathbf{K}}$ approximates \mathbf{K} well, which can be theoretically guaranteed, the approximate Laplacian $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{K}}\tilde{\mathbf{D}}^{-1/2}$ can be far from \mathbf{L} . This is because a small perturbation in $\tilde{\mathbf{D}}$ can have an out-sized influence on the eigenvectors of $\tilde{\mathbf{L}}$.
- One may propose to approximate $\mathbf{N} = \mathbf{D}^{-1/2}\mathbf{K}\mathbf{D}^{-1/2}$, rather than \mathbf{K} , with a Nyström approximation $\tilde{\mathbf{N}}$; this ensures that the approximate normalized graph Laplacian $\tilde{\mathbf{L}} = \mathbf{I}_n - \tilde{\mathbf{N}}$ is SPSD. However, this approach requires forming the entirety of \mathbf{K} in order to compute the degree matrix \mathbf{D} , and thus has quadratic (with n) time and memory costs. Furthermore, although the resulting approximation, $\tilde{\mathbf{L}}$, is SPSD, it is not a graph Laplacian: its off-diagonal entries are not guaranteed to be non-positive, and its smallest eigenvalue may be nonzero.

In summary, spectral clustering using the Nyström approximation (Fowlkes et al., 2004), which has proven to be a useful heuristic, and which is composed of theoretically principled parts, is less principled when viewed in its entirety. Approximate kernel k -means clustering using Nyström approximation is an equivalently efficient, but theoretically more principled alternative.

4.3. Empirical Comparison with Approximate Spectral Clustering using Nyström Approximation

To complement our discussion of the relative merits of the two methods, we empirically compared the performance of our novel method of approximate kernel k -means clustering

Table 3: Summary of the data sets used in our comparisons.

dataset	#instances (n)	#features (d)	#clusters (k)
MNIST (LeCun et al., 1998)	60,000	780	10
Mushrooms (Frank and Asuncion, 2010)	8,124	112	2
PenDigits (Frank and Asuncion, 2010)	7,494	16	10

using the Nyström method with the popular method of approximate spectral clustering using the Nyström method. We used three classification data sets, described in Table 3. The data sets used are available at <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.

Let $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$ be the input vectors. We take both the affinity matrix for spectral clustering and the kernel matrix for kernel k -means to be the RBF kernel matrix $\mathbf{K} = [\kappa(\mathbf{a}_i, \mathbf{a}_j)]_{ij} \in \mathbb{R}^{n \times n}$, where $\kappa(\mathbf{a}, \mathbf{a}') = \exp(-\frac{1}{2\sigma^2} \|\mathbf{a} - \mathbf{a}'\|_2^2)$ and σ is the kernel width parameter. We choose σ based on the average interpoint distance in the data sets as

$$\sigma = \beta \cdot \sqrt{\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \|\mathbf{a}_i - \mathbf{a}_j\|_2^2}, \quad (5)$$

where we take $\beta = 0.2, 1, \text{ or } 5$.

The algorithms under comparison are all implemented in Python 3.5.2. Our implementation of approximate spectral clustering follows the code in (Fowlkes et al., 2004). To compute linear k -means clusterings, we use the function `sklearn.cluster.KMeans` present in the scikit-learn package. Our algorithm for approximate kernel k -means clustering is described in more detail in Section 5.1. We ran the computations on a MacBook Pro with a 2.5GHz Intel Core i7 CPU and 16GB of RAM.



We compare approximate spectral clustering (SC) with approximate kernel k -means clustering (KK), with both using the rank-restricted Nyström method with uniform sampling.⁴ We used normalized mutual information (NMI) (Strehl and Ghosh, 2002) to evaluate clustering performance: the NMI falls between 0 (representing no mutual information between the true and approximate clusterings) and 1 (perfect correlation of the two clusterings), so larger NMI indicates better performance. The target dimension s is taken to be k ; and, for each method, the sketch size c is varied from $5s$ to $50s$. We record the time cost of the two methods, excluding the time spent on the k -means clustering required in both algorithms.⁵ We repeat this procedure 100 times and report the averaged NMI and average elapsed time.

We note that, at small sketch sizes c , exceptions often arise during approximate spectral clustering due to negative entries in the degree matrix. (This is an example, as discussed in Section 4.2, of when approximate spectral clustering heuristics do not perform well.) We discard the trials where such exceptions occur.



Our results are summarized in Figure 3. Figure 3 illustrates the NMI of SC and KK as a function of the sketch size c and as a function of elapsed time for both algorithms. While

4. Uniform sampling is appropriate for the value of σ used in Eqn. (5); see Gittens and Mahoney (2016) for a detailed discussion of the effect of varying σ .



5. For both SC and KK with Nyström approximation, the extracted feature matrices have dimension $n \times k$, so the k -means clusterings required by both SC and KK have identical cost.



(a) MNIST ($n = 60,000$, $d = 780$, $k = 10$).



(b) Mushroom ($n = 8,124$, $d = 112$, $k = 2$).



(c) PenDigits ($n = 7,494$, $d = 16$, $k = 10$).



Figure 3: Comparisons between the runtimes and accuracies of approximate spectral clustering (SC) and approximate kernel k -means clustering (KK), both using uniform column sampling Nyström approximation.

Algorithm 2 Approximate Kernel k -Means Clustering using Nyström Approximation.

-
- 1: **Input:** data set $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$, number of clusters k , target dimension $s (\geq k)$, arbitrary integer $\ell (> s)$, sketch size $c (> \ell)$, kernel function κ .
 - 2: // **Step 1: The Nyström Method**
 - 3: Form sketches $\mathbf{C} = \mathbf{K}\mathbf{P} \in \mathbb{R}^{n \times c}$ and $\mathbf{W} = \mathbf{P}^T \mathbf{C} \in \mathbb{R}^{c \times c}$, where $\mathbf{K} = [\kappa(\mathbf{a}_i, \mathbf{a}_j)]_{ij} \in \mathbb{R}^{n \times n}$ is the kernel matrix and $\mathbf{P} \in \mathbb{R}^{n \times c}$ is some sketching matrix, e.g., uniform sampling;
 - 4: Compute a matrix $\mathbf{R} \in \mathbb{R}^{n \times \ell}$ such that $\mathbf{R}\mathbf{R}^T = \mathbf{C}\mathbf{W}_\ell^{-1}\mathbf{C}^T$;
 - 5: // **Step 2: Dimensionality Reduction**
 - 6: Compute the rank- s truncated SVD $\mathbf{R}_s = \tilde{\mathbf{U}}_s \tilde{\Sigma}_s \tilde{\mathbf{V}}_s^T$;
 - 7: Let $\mathbf{B} = \tilde{\mathbf{U}}_s \tilde{\Sigma}_s (= \mathbf{R}\tilde{\mathbf{V}}_s) \in \mathbb{R}^{n \times s}$;
 - 8: // **Step 3: Linear k -Means Clustering**
 - 9: Perform k -means clustering over the rows of \mathbf{B} ;
 - 10: **return** the clustering results.
-

there are quantitative differences between the results on the three data sets, the plots all show that KK is more accurate as a function of the sketch size or elapsed time than SC.

5. Single-Machine Medium-Scale Experiments

In this section, we empirically compare the Nyström method and random feature maps (Rahimi and Recht, 2007) for kernel k -means clustering. We conduct experiments on the data listed in Table 3. For the Mushrooms and PenDigits data, we are able to evaluate the objective function value of kernel k -means clustering.

5.1. Single-Machine Implementation of Approximate Kernel k -Means

Our algorithm for approximate kernel k -means clustering comprises three steps: **Nyström approximation**, **dimensionality reduction**, and **linear k -means clustering**. Both the single-machine as well as the distributed variants of the algorithm are governed by three parameters: s , the number of features used in the clustering; ℓ , a regularization parameter; and c , the sketch size. These parameters satisfy $k \leq s < \ell \leq c \ll n$.

1. **Nyström approximation.** Let c be the sketch size and $\mathbf{P} \in \mathbb{R}^{n \times c}$ be a sketching matrix. Let $\mathbf{C} = \mathbf{K}\mathbf{P}$ and $\mathbf{W} = \mathbf{P}^T \mathbf{K} \mathbf{P} = \mathbf{P}^T \mathbf{C}$. The standard Nyström approximation is $\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T$; small singular values in \mathbf{W} can lead to instability in the Moore-Penrose inverse, so a widely used heuristic is to choose $\ell < c$ and use $\mathbf{C}\mathbf{W}_\ell^{-1} \mathbf{C}^T$ instead of the standard Nyström approximation.⁶ We set $\ell = \lceil c/2 \rceil$ (arbitrarily). Let $\mathbf{W}_\ell = \mathbf{U}_{\mathbf{W},\ell} \mathbf{\Lambda}_{\mathbf{W},\ell} \mathbf{U}_{\mathbf{W},\ell}^T$ be the truncated SVD of \mathbf{W} and return $\mathbf{R} = \mathbf{C}\mathbf{U}_{\mathbf{W},\ell} \mathbf{\Lambda}_{\mathbf{W},\ell}^{-1/2} \in \mathbb{R}^{n \times \ell}$ as the output of the Nyström method.

6. The Nyström approximation $\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T$ is correct in theory, but the Moore-Penrose inverse often causes numerical errors in practice. The Moore-Penrose inverse drops all the zero singular values, however, due to the finite numerical precision, it is difficult to determine whether a singular value, say 10^{-12} , should be zero or not, and this makes the computation unstable: if such a small singular value is believed to be zero, it will be dropped; otherwise, the Moore-Penrose inverse will invert it to obtain a singular value of 10^{12} . Dropping some portion of the smallest singular values is a simple heuristic that avoids this instability. This is why we heuristically use $\mathbf{C}\mathbf{W}_\ell^{-1} \mathbf{C}^T$ instead of $\mathbf{C}\mathbf{W}^\dagger \mathbf{C}^T$. Currently we do not have theory for this heuristic. Chiu and Demanet (2013) considers the theoretical implications of this regularization heuristic, but their results do not apply to our problem.

2. **Dimensionality reduction.** Let $\tilde{\mathbf{V}}_s \in \mathbb{R}^{\ell \times s}$ contain the dominant s right singular vectors of \mathbf{R} . Let $\mathbf{B} = \mathbf{R}\tilde{\mathbf{V}}_s \in \mathbb{R}^{n \times s}$. It can be verified that $\mathbf{B}\mathbf{B}^T = (\mathbf{C}\mathbf{W}_\ell^{-1}\mathbf{C}^T)_s$, which is our desired rank-restricted Nyström approximation.
3. **Linear k -means clustering.** With $\mathbf{B} \in \mathbb{R}^{n \times s}$ at hand, use an arbitrary off-the-shelf linear k -means clustering algorithm to cluster the rows of \mathbf{B} .

See Algorithm 2 for the single-machine version of this approximate kernel k -means clustering algorithm. Observe that we can use uniform sampling to form \mathbf{C} and \mathbf{W} , and thereby avoid computing most of \mathbf{K} .

Let $\mathbf{R} \in \mathbb{R}^{n \times c}$ be the feature matrix computed by random feature maps (RFM). To make the comparison fair, we perform **dimensionality reduction** for RFM (Rahimi and Recht, 2007) in the same way as described in Algorithm 2 to compute $\mathbf{B} = \mathbf{R}\tilde{\mathbf{V}}_s$, and apply linear k -means clustering on the rows of \mathbf{B} .

5.2. Comparing Nyström, Random Feature Maps, and Two-Step Method

We empirically compare the clustering performances of kernel approximations formed using Nyström, random feature map (RFM) (Rahimi and Recht, 2007), and the two-step method (Chitta et al., 2011) on the data sets detailed in Table 3.

We use the RBF kernel with width parameter given by (5); Figure 3 indicates that $\beta = 1$ is a good choice for these data sets. We conduct dimensionality reduction for both Nyström and RFM to obtain s -dimensional features, and consider three choices: $s = k$, $s = \lceil \sqrt{ck} \rceil$, and without dimensionality reduction (equivalently, $s = c$).

The quality of the clusterings is quantified using both normalized mutual information (NMI) (Strehl and Ghosh, 2002) and the objective function value:

$$\frac{1}{n} \sum_{i=1}^k \sum_{j \in \mathcal{J}_i} \left\| \mathbf{k}_j - \frac{1}{|\mathcal{J}_i|} \sum_{l \in \mathcal{J}_i} \mathbf{k}_l \right\|_2^2, \quad (6)$$


where $\mathbf{k}_1, \dots, \mathbf{k}_n \in \mathbb{R}^n$ are the columns of the kernel matrix \mathbf{K} , and the disjoint sets $\mathcal{J}_1, \dots, \mathcal{J}_k$ reflect the clustering.

We repeat the experiments 500 times and report the results in Figures 4 and 5. The experiments show that as measured by both NMIs and objective values, the Nyström method outperforms RFM in most cases. Both the Nyström method and RFM are consistently superior to the two-step method of (Chitta et al., 2011), which requires a large sketch size. All the compared methods improve as the sketch size c increases.


Judging from these medium-scale experiments, the target rank s has little impact on the NMI and clustering objective value. This phenomenon is not general; in the large-scale experiments of the next section we see that setting s properly allows one to obtain a better NMI than an over-small or over-large s .

6. Large-Scale Experiments using Distributed Computing


In this section, we empirically study our approximate kernel k -means clustering algorithm on large-scale data. We state a distributed version of the algorithm, implement it in Apache



(a) MNIST ($n = 60,000$, $d = 780$, $k = 10$). The results of Nystrom and RFM are plotted in the figures. The medians of the two-step method with $c = 3k$, $9k$, $27k$, and $81k$ are 0.072, 0.012, 0.017, and 0.032, respectively.




(b) Mushroom ($n = 8,124$, $d = 112$, $k = 2$). The results of Nystrom and RFM are plotted in the figures. The medians of the two-step method with $c = 3k$, $9k$, $27k$, and $81k$ are 0.123, 0.224, 0.263, and 0.494, respectively.




(c) PenDigits ($n = 7,494$, $d = 16$, $k = 10$). The results of Nystrom and RFM are plotted in the figures. The medians of the two-step method with $c = 3k$, $9k$, $27k$, and $81k$ are 0.399, 0.413, 0.422, and 0.421, respectively.

— Nystrom — Random Feature

Figure 4: Quality of approximate kernel k -means clusterings using Nystrom, random feature maps (Rahimi and Recht, 2007), and the two-step method (Chitta et al., 2011). (Dimensionality reduction is not applicable to the two-step method). The y -axis reports the normalized mutual information (NMI).



(a) Mushroom ($n = 8,124$, $d = 112$, $k = 2$). The results of Nyström and RFM are plotted in the figures. The medians of the objective values of the two-step method with $c = 3k$, $9k$, $27k$, and $81k$ are 101.5, 92.8, 89.5, and 85.3, respectively. A random partition results in a median objective function value of 119.5.



(b) PenDigits ($n = 7,494$, $d = 16$, $k = 10$). The results of Nyström and RFM are plotted in the figures. The medians of the two-step method with $c = 3k$, $9k$, $27k$, and $81k$ are 210.5, 198.2, 193.5, and 193.4, respectively. A random partition results in a median objective function value of 395.4.

— Nystrom — Random Feature

Figure 5: Quality of approximate kernel k -means clusterings obtained using rank-restricted Nyström, RFMs (Rahimi and Recht, 2007), and the two-step method (Chitta et al., 2011). (Dimensionality reduction is not applicable to the two-step method). The y -axis reports the kernel k -means clustering objective function value.

Algorithm 3 Distributed Approximate Kernel k -Means Clustering using Nyström Approximation.

- 1: **Input:** data set $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$ distributed among m machines, number of clusters k , target dimension s ($\geq k$), arbitrary integer $\ell > s$, sketch size c ($> \ell$), kernel function κ .
 - 2: // **Step 1: The Nyström Method**
 - 3: Sample c vectors from $\mathbf{a}_1, \dots, \mathbf{a}_n$ to form $\mathbf{a}'_1, \dots, \mathbf{a}'_c$ and send to the driver;
 - 4: Driver computes $\mathbf{W} = [\kappa(a'_i, a'_j)]_{ij} \in \mathbb{R}^{c \times c}$ and a matrix $\mathbf{Z} \in \mathbb{R}^{c \times \ell}$ satisfying $\mathbf{Z}\mathbf{Z}^T = \mathbf{W}_\ell^{-1}$;
 - 5: Broadcast \mathbf{Z} and $\mathbf{a}'_1, \dots, \mathbf{a}'_c$ to all executors;
 - 6: **for** all $j \in [n]$ **do**
 - 7: Each executor locally computes $\mathbf{c}_j = [\kappa(\mathbf{a}_j, \mathbf{a}'_1); \dots; \kappa(\mathbf{a}_j, \mathbf{a}'_c)] \in \mathbb{R}^c$ and $\mathbf{r}_j = \mathbf{Z}^T \mathbf{c}_j \in \mathbb{R}^\ell$;
 - 8: **end for**
 - 9: // **Step 2: Dimensionality Reduction**
 - 10: Let $\tilde{\mathbf{V}}_s$ contain the dominant s right singular vectors of the distributed row matrix $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_n]^T \in \mathbb{R}^{n \times \ell}$ (computed using a distributed truncated SVD algorithm);
 - 11: Broadcast $\tilde{\mathbf{V}}_s \in \mathbb{R}^{\ell \times s}$ to all executors;
 - 12: **for** all $j \in [n]$ **do**
 - 13: Each executor locally computes $\mathbf{b}_j = \tilde{\mathbf{V}}_s^T \mathbf{r}_j \in \mathbb{R}^s$;
 - 14: **end for**
 - 15: // **Step 3: Linear k -Means Clustering**
 - 16: Perform k -means clustering over the Nyström features $\mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^s$ (using a distributed linear k -means clustering algorithm).
 - 17: **return** the clustering results.
-

Spark⁷, and evaluate its performance on NERSC’s Cori supercomputer. We investigate the effect of increased parallelism, sketch size c , and target dimension s .

Algorithm 3 is a distributed version of our method described in Section 5.1. Again, we use uniform sampling to form \mathbf{C} and \mathbf{W} to avoid computing most of \mathbf{K} . We mainly focus on the **Nyström approximation** step, as the other two steps are well supported by distributed computing systems such as Apache Spark.

6.1. Experimental Setup

We implemented Algorithm 3 in the Apache Spark framework (Zaharia et al., 2010, 2012), using the Scala API. We computed the Nyström approximation using the matrix operations provided by Spark, and invoked the MLlib library for machine learning in Spark (Meng et al., 2016) to perform the dimensionality reduction and linear k -means clustering steps. For the linear k -means clustering, we set the maximum number of iterations to 100.

We ran our experiments on Cori Phase I, a NERSC supercomputer, located at Lawrence Berkeley National Laboratory. Cori Phase I is a Cray XC40 system with 1632 compute nodes, each of which has two 2.3GHz 16-core Haswell processors and 128GB of DRAM. The Cray Aries high-speed interconnect linking the compute nodes is configured in a dragonfly topology.

We used the MNIST8M data set to conduct our empirical evaluations; this data set has $n = 8.1 \times 10^6$ instances, $d = 784$ features, and $k = 10$ clusters. We vary c and s and set $\ell = \frac{c}{2}$. We chose the RBF kernel width parameter according to (5) with $\beta = 1.0$. We use

7. This implementation is available at <https://github.com/wangshusen/SparkKernelKMeans.git>.

Table 4: The mean and standard deviation of the runtimes of Algorithm 3, in seconds, as a function of the number of compute nodes. We report the total runtime as well as the runtimes of the three stages: **Nyström approximation**, **dimensionality reduction (DR)**, and **linear k -means clustering**. The total time is the sum of the three stages and all the Spark overheads, e.g., Spark initialization. We fix $k = 10$ and $c = 400$. In the upper table, we set $s = 20$, and the NMI is 0.400 ± 0.009 ; in the lower table, we set $s = 80$, and the NMI is 0.410 ± 0.010 .

	8 Nodes	16 Nodes	32 Nodes	64 Nodes	128 Nodes
Nyström	2998.6 \pm 263.9	1268.8 \pm 106.4	665.2 \pm 39.6	369.7 \pm 126.1	181.4 \pm 22.2
DR	37.4 \pm 58.3	63.8 \pm 72.5	38.9 \pm 18.2	87.3 \pm 37.8	183.9 \pm 89.1
k-means	117.2 \pm 132.3	153.1 \pm 98.8	119.9 \pm 76.9	223.2 \pm 88.4	391.0 \pm 155.9
Total	3201.5 \pm 344.0	1532.0 \pm 146.9	867.9 \pm 97.5	734.7 \pm 152.0	828.2 \pm 210.4

	8 Nodes	16 Nodes	32 Nodes	64 Nodes	128 Nodes
Nyström	3008.4 \pm 385.6	1312.9 \pm 141.2	696.0 \pm 111.8	342.5 \pm 16.9	197.0 \pm 24.1
DR	53.0 \pm 41.9	58.7 \pm 22.7	94.6 \pm 31.8	179.4 \pm 55.8	470.6 \pm 118.5
k-means	58.2 \pm 24.7	80.9 \pm 31.8	104.6 \pm 28.7	211.3 \pm 103.1	501.7 \pm 162.9
Total	3168.1 \pm 434.1	1492.6 \pm 135.3	940.6 \pm 139.3	775.9 \pm 151.6	1232.4 \pm 232.1

Cori’s default setting of Spark configurations. For each setting of parameters, we repeated the experiments 10 times and recorded the NMIs and elapsed time.

6.2. Effect of Increased Parallelism

We varied the number of nodes to test the impact of increased parallelism on each of the three steps in Algorithm 3. We set the sketch size to $c = 400$ and the target dimension to $s = 20$ or 80 . Table 4 reports the mean and standard deviation of the elapsed times. As a reference point, using 32 nodes, our algorithm takes 15 minutes on average to group the 8.1 million input instances into 10 clusters. We also plot the elapsed times in Figure 6.

The **Nyström method** scales very well with the increase of nodes: its elapsed time is inversely proportional to the number of nodes. This is because the **Nyström method** requires only 2 rounds of communications; the elapsed time is spent mostly on computation on the executors.

Dimensionality reduction (DR) and **linear k -means clustering** are highly iterative and thus have high latency, incur straggler delays, and have large scheduler and communication overheads. See (Gittens et al., 2016) for an in-depth discussion of the performance concerns when using Spark for distributed matrix computations. As the number of nodes increases from 32 to 64 and higher, **DR** and **k -means** exhibit anti-scaling behavior: as the number of nodes goes up, their runtimes increase rather than decrease. For **DR**, the input dimension is $\ell = \frac{c}{2} = 200$, and the target dimension is $s = 20$ or 80 ; while for **k -means**, the input dimension is $s = 20$ or 80 . Clearly, the issue here is not that **DR** and **k -means** are computationally intensive. Instead, as the number of nodes increase, although the per-node computational time decreases, the Spark communications overheads are increasing, making **DR** and **k -means** less performant. Such




Figure 6: The mean of the runtimes of Algorithm 3, in seconds, as a function of the number of compute nodes.

behavior has been characterized and studied in other Spark implementations of linear algebra algorithms (Gittens et al., 2016).

Table 5: The NMI and elapsed time (seconds) for varying sketch sizes c . Here “T” denotes the elapsed time. We use 32 nodes. In the upper table, we set $s = 20$; in the lower table, we set $s = 80$.

	$c = 100$	$c = 400$	$c = 1,600$
NMI	0.3833 ± 0.0117	0.3975 ± 0.0112	0.4069 ± 0.0001
T(Nyström)	138.5 ± 62.1	665.2 ± 39.5	2634.5 ± 94.4
T(DR)	1.2 ± 0.5	38.9 ± 18.2	31.3 ± 6.6
T(k -Means)	60.9 ± 28.7	119.9 ± 76.9	86.9 ± 27.3
T(Total)	227.4 ± 101.3	867.9 ± 97.5	2785.4 ± 83.2

	$c = 400$	$c = 1600$
NMI	0.4101 ± 0.0101	0.4233 ± 0.0131
T(Nyström)	696.0 ± 111.8	2728.1 ± 190.3
T(DR)	94.6 ± 31.8	97.9 ± 20.6
T(k -Means)	104.6 ± 28.7	88.0 ± 23.9
T(Total)	940.6 ± 139.3	2952.2 ± 174.5

6.3. Effect of Sketch Size c

We executed our Spark implementation using 32 compute nodes, setting $s = 20$ or 80 , and varying the target dimension c . Table 5 reports the observed normalized mutual information




Figure 7: The mean of the runtimes of Algorithm 3, in seconds, as a function of the sketch size c .

(NMI) (Strehl and Ghosh, 2002) and elapsed times. As predicted by our theory, for fixed k and s , larger c always leads to better performance.

In Figure 7 we plot the runtime as a function of the sketch size c . The plots shows the significant influence of c on the running time. In Figure 7, the runtime of **Nyström** grows superlinearly with c . According to our analysis, the computational and communication costs of **Nyström** are both superlinear in c . Therefore, the user should keep in mind this trade-off between the clustering performance and the computational cost.

6.4. Effect of Target Dimension s

We executed our Spark implementation using 32 computational nodes, fixing $k = 10$ and $c = 1600$ and varying the target dimension s . Table 6 reports the observed NMIs and elapsed times. For fixed k and c , a moderately large s leads to better clustering performance. However, as s grows, $\frac{k}{s}$ decreases but $\frac{s}{c}$ increases, so as predicted in Remark 3, the clustering performance is nonmonotonic in s . Indeed, as s grows from 80 to 160, the NMI deteriorates. In practice, for fixed k and c , one should set s moderately large, but not over-large.

In Figure 8 we plot the runtime as a function of the target dimension s . Note that s does not affect the time cost of **the Nyström method**. As s grows, the time costs of **dimensionality reduction** and **linear k -means clustering** both increase. This implies that a moderate s is good for computational purpose. Previously, in Figure 2, we plot the NMI as a function of s while fixing k and c . Figure 2 shows that a proper s leads to better NMI than an over-large or over-small s . In sum, setting s properly has both computational and accuracy benefits, which corroborates our theory.

Table 6: The NMI and elapsed time (seconds) for varying target dimensions s . Here “T” denotes the elapsed time. We use 32 nodes and fix $c = 1600$.

	$s = 10$	$s = 20$	$s = 40$	$s = 80$	$s = 160$	$s = 320$
NMI	0.3852 ± 0.0003	0.4069 ± 0.0001	0.4130 ± 0.0062	0.4233 ± 0.0131	0.4086 ± 0.0125	0.4099 ± 0.0098
T(Nyström)	2702.2 ± 131.6	2634.5 ± 94.4	2815.3 ± 403.6	2728.1 ± 190.3	2849.9 ± 511.2	2757.1 ± 271.9
T(DR)	21.1 ± 3.2	31.3 ± 6.6	57.5 ± 9.2	97.9 ± 20.6	168.2 ± 29.0	570.5 ± 151.0
T(k -Means)	84.0 ± 6.4	86.9 ± 27.3	90.6 ± 9.0	88.0 ± 23.9	99.8 ± 12.2	169.3 ± 64.2
T(Total)	2840.9 ± 132.0	2785.4 ± 83.2	3003.4 ± 413.7	2952.2 ± 174.5	3157.5 ± 549.2	3537.5 ± 256.3




Figure 8: The mean of the runtimes of Algorithm 3, in seconds, as a function of the target dimension s . We use 32 nodes and fix $k = 10$ and $c = 1600$.

7. Conclusion

We provided principled algorithms for computing approximate kernel k -means clusterings. In particular, we showed that the combination of linear k -means with rank-restricted Nyström approximation is theoretically sound, practically useful, and scalable to large data sets. This should be contrasted with approximate spectral clustering using Nyström approximation. Although the latter is a widely-used approach to scalable non-linear clustering, it has theoretical deficiencies and practical limitations. Experiments demonstrated that approximate kernel k -means clustering using the rank-restricted Nyström approximation consistently outperforms approximate spectral clustering.

Our analysis uses the concept of projection-cost preservation and builds upon the existing theory of randomized linear algebra. Our main result is a $1 + \epsilon$ approximation ratio guarantee for kernel k -means clustering: when $s = \frac{k}{\epsilon}$ Nyström features are used, a $1 + \mathcal{O}(\epsilon)$ approximation ratio is guaranteed with high probability. As an intermediate theoretical result of independent interest, we introduced a novel rank-restricted Nyström approximation and proved that it gives a $1 + \epsilon$ relative-error low-rank approximation guarantee in the trace norm.

To complement our theory, we investigated the performance of the rank-restricted Nyström approximation when applied to kernel k -means clustering. We produced an Apache Spark implementation of a distributed version of our approximate kernel k -means algorithm, and applied it to cluster 8.1 million vectors; the results demonstrate the usefulness and simplicity of kernel k -means with rank-restricted Nyström approximation for clustering moderately large-scale data sets. Combined with recent work on user-friendly frameworks for distributed computing systems (Zaharia et al., 2010, 2012), large-scale machine learning (Meng et al., 2016), and large-scale randomized linear algebra (Gittens et al., 2016), our results suggest the use of kernel k -means with rank-restricted Nyström approximations for simple, computationally efficient, and theoretically principled clustering of large-scale data sets.

Acknowledgments

We thank the authors of Musco and Musco (2017) for bringing to our attention their results on approximate kernel k -means clustering through the use of ridge leverage score Nyström approximations, the authors of Tropp et al. (2017) for informing us of their contemporaneous guarantees on the approximation error of rank-restricted Nyström approximations, and the anonymous reviewers for their helpful suggestions. We would like to acknowledge ARO, DARPA, NSF, and ONR for providing partial support of this work.

Appendix A. Proof of Theorem 1

In this section, we will provide a proof of Theorem 1, our main quality-of-approximation result for rank-restricted Nyström approximation. We will start, in Section A.1, by establishing three technical lemmas; then, in Section A.2, we will establish an important structural result on the rank-restricted Nyström approximation as well as key approximate matrix multiplication properties for low-rank matrix approximation; and finally, in Section A.3, we will use these results to prove Theorem 1.

A.1. Technical Lemmas

Lemma 5 is a very well known result. See (Boutsidis et al., 2014; Zhang, 2015). Here we offer a simplified proof.

Lemma 5 *Let $\mathbf{A} \in \mathbb{R}^{n \times m}$ be any matrix and $\mathbf{U} \in \mathbb{R}^{n \times c}$ have orthonormal columns. Let s be any positive integer no greater than c . Then*

$$(\mathbf{U}^T \mathbf{A})_s = \operatorname{argmin}_{\operatorname{rank}(\mathbf{Z}) \leq s} \|\mathbf{A} - \mathbf{UZ}\|_F^2.$$

Proof Let $\mathbf{U}^\perp \in \mathbb{R}^{n \times (n-c)}$ be the orthogonal complement of \mathbf{U} . It holds that

$$\begin{aligned}
 \|\mathbf{A} - \mathbf{UZ}\|_F^2 &= \left\| \mathbf{A} - \begin{bmatrix} \mathbf{U} & \mathbf{U}^\perp \end{bmatrix} \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} \right\|_F^2 \\
 &= \left\| \begin{bmatrix} \mathbf{U}^T \\ (\mathbf{U}^\perp)^T \end{bmatrix} \mathbf{A} - \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} \right\|_F^2 \\
 &= \left\| \begin{bmatrix} \mathbf{U}^T \mathbf{A} - \mathbf{Z} \\ (\mathbf{U}^\perp)^T \mathbf{A} \end{bmatrix} \right\|_F^2 \\
 &= \|\mathbf{U}^T \mathbf{A} - \mathbf{Z}\|_F^2 + \|(\mathbf{U}^\perp)^T \mathbf{A}\|_F^2,
 \end{aligned}$$

where the second equality follows from the unitary invariance of the Frobenius norm. Then

$$\operatorname{argmin}_{\operatorname{rank}(\mathbf{Z}) \leq s} \|\mathbf{A} - \mathbf{UZ}\|_F^2 = \operatorname{argmin}_{\operatorname{rank}(\mathbf{Z}) \leq s} \|\mathbf{U}^T \mathbf{A} - \mathbf{Z}\|_F^2 = (\mathbf{U}^T \mathbf{A})_s,$$

by which the lemma follows. ■

Lemma 6 is a new result established in this work. The lemma plays an important role in analyzing the rank-restricted Nystrom method.

Lemma 6 *Let $\mathbf{A} \in \mathbb{R}^{n \times m}$ be any matrix and $\mathbf{U} \in \mathbb{R}^{n \times c}$ has orthonormal columns. Let $s \leq c$ be arbitrary integer. Let $\mathbf{Q} \in \mathbb{R}^{n \times s}$ be the orthonormal bases of the rank s matrix $\mathbf{U}(\mathbf{U}^T \mathbf{A})_s \in \mathbb{R}^{n \times m}$. Then $\mathbf{Q}\mathbf{Q}^T \mathbf{U}(\mathbf{U}^T \mathbf{A})_s = \mathbf{U}(\mathbf{U}^T \mathbf{A})_s$.*

Proof Let $\mathbf{Z}^* = \mathbf{Q}^T \mathbf{U}(\mathbf{U}^T \mathbf{A})_s \in \mathbb{R}^{s \times m}$. By the definition of \mathbf{Q} , we have

$$\mathbf{Q}\mathbf{Q}^T \mathbf{U}(\mathbf{U}^T \mathbf{A})_s = \mathbf{U}(\mathbf{U}^T \mathbf{A})_s. \tag{7}$$

It holds that

$$\|\mathbf{U}(\mathbf{U}^T \mathbf{A})_s - \mathbf{A}\|_F^2 = \min_{\operatorname{rank}(\mathbf{X}) \leq s} \|\mathbf{UX} - \mathbf{A}\|_F^2 \leq \min_{\operatorname{rank}(\mathbf{Z}) \leq s} \|\mathbf{QZ} - \mathbf{A}\|_F^2 \leq \|\mathbf{QZ}^* - \mathbf{A}\|_F^2,$$

where the equality follows from Lemma 5; the former inequality follows from that $\operatorname{range}(\mathbf{Q}) \subset \operatorname{range}(\mathbf{U})$. Because $\mathbf{U}(\mathbf{U}^T \mathbf{A})_s = \mathbf{QZ}^*$ by (7), the lefthand and righthand sides are equal, and thereby

$$\min_{\operatorname{rank}(\mathbf{Z}) \leq s} \|\mathbf{QZ} - \mathbf{A}\|_F^2 = \|\mathbf{QZ}^* - \mathbf{A}\|_F^2.$$

Since \mathbf{Z} is $s \times m$, it follows that

$$\min_{\mathbf{Z}} \|\mathbf{QZ} - \mathbf{A}\|_F^2 = \min_{\operatorname{rank}(\mathbf{Z}) \leq s} \|\mathbf{QZ} - \mathbf{A}\|_F^2 = \|\mathbf{QZ}^* - \mathbf{A}\|_F^2. \tag{8}$$

Let $\mathbf{Q}^\perp \in \mathbb{R}^{n \times (n-s)}$ be the orthogonal complement of \mathbf{Q} . It holds that

$$\begin{aligned} \|\mathbf{A} - \mathbf{QZ}\|_F^2 &= \left\| \mathbf{A} - \begin{bmatrix} \mathbf{Q} & \mathbf{Q}^\perp \end{bmatrix} \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} \right\|_F^2 \\ &= \left\| \begin{bmatrix} \mathbf{Q}^T \\ (\mathbf{Q}^\perp)^T \end{bmatrix} \mathbf{A} - \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} \right\|_F^2 \\ &= \left\| \begin{bmatrix} \mathbf{Q}^T \mathbf{A} - \mathbf{Z} \\ (\mathbf{Q}^\perp)^T \mathbf{A} \end{bmatrix} \right\|_F^2 \\ &= \|\mathbf{Q}^T \mathbf{A} - \mathbf{Z}\|_F^2 + \|(\mathbf{Q}^\perp)^T \mathbf{A}\|_F^2, \end{aligned}$$

where the second equality follows from the unitary invariance of the Frobenius norm. Obviously, $\mathbf{Z} = \mathbf{Q}^T \mathbf{A}$ is the unique minimizer of $\min_{\mathbf{Z}} \|\mathbf{QZ} - \mathbf{A}\|_F^2$; otherwise $\mathbf{Q}^T \mathbf{A} - \mathbf{Z}$ in the righthand side is non-zero, making the objective function increase.

On the one hand, Eqn. (8) shows that $\mathbf{Z}^* = \mathbf{Q}^T \mathbf{U}(\mathbf{U}^T \mathbf{A})_s \in \mathbb{R}^{s \times m}$ is one minimizer of $\min_{\mathbf{Z}} \|\mathbf{QZ} - \mathbf{A}\|_F^2$. On the other hand, we have that $\mathbf{Z} = \mathbf{Q}^T \mathbf{A}$ is the unique minimizer of $\min_{\mathbf{Z}} \|\mathbf{QZ} - \mathbf{A}\|_F^2$. By the uniqueness, we have

$$\mathbf{Z}^* = \mathbf{Q}^T \mathbf{U}(\mathbf{U}^T \mathbf{A})_s = \mathbf{Q}^T \mathbf{A}.$$

It follows that

$$\mathbf{Q} \mathbf{Q}^T \mathbf{U}(\mathbf{U}^T \mathbf{A})_s = \mathbf{Q} \mathbf{Q}^T \mathbf{A}.$$

The lemma follows from (7) and the above equality. \blacksquare

Lemma 7 is a new result established by this work. Let \mathbf{A} be any symmetric matrix. Lemma 7 analyzes the power scheme: using $\mathbf{A}^t \mathbf{P}$, instead of $\mathbf{A} \mathbf{P}$, as a sketch of \mathbf{A} . The lemma shows that the power scheme leads to an improvement of $(\frac{\sigma_{s+1}^2}{\sigma_s^2})^{t-1}$, where σ_i is the i -th biggest singular value of \mathbf{A} . Lemma 7 with $t = 1$ is identical to (Boutsidis et al., 2014, Lemma 9). The power scheme has been studied by Gittens and Mahoney (2016); Halko et al. (2011); Woodruff (2014); their results do not allow the rank restriction $\text{rank}(\mathbf{X}) \leq s$. To prove Theorem 1, we only need $t = 1$; we will use Lemma 7 to extend Theorem 1 to the power method.

Lemma 7 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be any SPSD matrix and $\mathbf{A}_s = \mathbf{V}_s \boldsymbol{\Sigma}_s \mathbf{V}_s^T$ be the truncated SVD. Let $\mathbf{P} \in \mathbb{R}^{n \times c}$ satisfy that $\mathbf{P}^T \mathbf{V}_s \in \mathbb{R}^{c \times s}$ has full column rank. Let $t \geq 1$ be any integer and $\mathbf{C} = \mathbf{A}^t \mathbf{P}$. Then for $\xi = 2$ or F ,*

$$\min_{\text{rank}(\mathbf{X}) \leq s} \|\mathbf{A} - \mathbf{CX}\|_\xi^2 \leq \|\mathbf{A} - \mathbf{A}_s\|_\xi^2 + \left(\frac{\sigma_{s+1}^2}{\sigma_s^2}\right)^{t-1} \|(\mathbf{A} - \mathbf{A}_s) \mathbf{P} (\mathbf{V}_s^T \mathbf{P})^\dagger\|_\xi^2,$$

where σ_i is the i -th largest singular value of \mathbf{A} .

Proof We construct the rank s matrix $\tilde{\mathbf{A}}_s = \mathbf{C} (\mathbf{V}_s^T \mathbf{P})^\dagger \boldsymbol{\Sigma}_s^{1-t} \mathbf{V}_s^T$ and use it to facilitate our proof. Since $\tilde{\mathbf{A}}_s$ has rank s , and its column space is in $\text{range}(\mathbf{C})$, it holds that

$$\begin{aligned} \min_{\text{rank}(\mathbf{X}) \leq s} \|\mathbf{A} - \mathbf{CX}\|_\xi^2 &\leq \|\mathbf{A} - \tilde{\mathbf{A}}_s\|_\xi^2 \\ &= \|(\mathbf{A} - \mathbf{A}_s) + (\mathbf{A}_s - \tilde{\mathbf{A}}_s)\|_\xi^2 \leq \|\mathbf{A} - \mathbf{A}_s\|_\xi^2 + \|\mathbf{A}_s - \tilde{\mathbf{A}}_s\|_\xi^2. \end{aligned}$$

The latter inequality follows from that the row space of $\mathbf{A}_s - \tilde{\mathbf{A}}_s$ is in $\text{range}(\mathbf{V}_s)$, that $(\mathbf{A} - \mathbf{A}_s)\mathbf{V}_s = \mathbf{0}$, and the matrix Pythagorean theorem. By the assumption $\text{rank}(\mathbf{V}_s^T \mathbf{P}) = s$, it holds that $(\mathbf{V}_s^T \mathbf{P})(\mathbf{V}_s^T \mathbf{P})^\dagger = \mathbf{I}_s$. We have

$$\begin{aligned}
 \|\mathbf{A}_s - \tilde{\mathbf{A}}_s\|_\xi^2 &= \|\mathbf{A}_s - (\mathbf{A}_s^t + \mathbf{A}^t - \mathbf{A}_s^t)\mathbf{P}(\mathbf{V}_s^T \mathbf{P})^\dagger \Sigma_s^{1-t} \mathbf{V}_s^T\|_\xi^2 \\
 &= \|\mathbf{A}_s - \underbrace{\mathbf{V}_s \Sigma_s^t (\mathbf{V}_s^T \mathbf{P})}_{s \times c} \underbrace{(\mathbf{V}_s^T \mathbf{P})^\dagger}_{c \times s} \Sigma_s^{1-t} \mathbf{V}_s^T - (\mathbf{A}^t - \mathbf{A}_s^t)\mathbf{P}(\mathbf{V}_s^T \mathbf{P})^\dagger \Sigma_s^{1-t} \mathbf{V}_s^T\|_\xi^2 \\
 &= \|\mathbf{A}_s - \mathbf{V}_s \Sigma_s \mathbf{V}_s^T - (\mathbf{A}^t - \mathbf{A}_s^t)\mathbf{P}(\mathbf{V}_s^T \mathbf{P})^\dagger \Sigma_s^{1-t} \mathbf{V}_s^T\|_\xi^2 \\
 &= \|(\mathbf{A}^t - \mathbf{A}_s^t)\mathbf{P}(\mathbf{V}_s^T \mathbf{P})^\dagger \Sigma_s^{1-t} \mathbf{V}_s^T\|_\xi^2 \\
 &\leq \|(\mathbf{A} - \mathbf{A}_s)^{t-1}\|_2^2 \|\Sigma_s^{1-t}\|_2^2 \|(\mathbf{A} - \mathbf{A}_s)\mathbf{P}(\mathbf{V}_s^T \mathbf{P})^\dagger\|_\xi^2 \\
 &= \left(\frac{\sigma_{s+1}^2}{\sigma_s^2}\right)^{t-1} \|(\mathbf{A} - \mathbf{A}_s)\mathbf{P}(\mathbf{V}_s^T \mathbf{P})^\dagger\|_\xi^2,
 \end{aligned}$$

by which the lemma follows. \blacksquare

A.2. Structural Result and Approximate Matrix Multiplication Properties

Lemma 8 is a structural result of independent interest that is important for establishing Theorem 1. By *structural result*, we mean that it is a linear algebraic result that holds for any matrix \mathbf{P} , and in particular for any matrix \mathbf{P} that is a randomized sketching matrix. In particular, given this structural result, randomness in an algorithm enters only via \mathbf{P} . The importance of establishing such structural results within randomized linear algebra has been highlighted previously (Drineas et al., 2011; Mahoney, 2011; Gittens and Mahoney, 2016; Mahoney and Drineas, 2016; Wang et al., 2017).

Lemma 8 *Let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be any SPSD matrix, $\mathbf{P} \in \mathbb{R}^{n \times c}$ be any matrix, and $s \leq c$ be any positive integer. Let $\mathbf{D} = \mathbf{K}^{1/2} \mathbf{P} \in \mathbb{R}^{n \times c}$, $\rho = \text{rank}(\mathbf{D}) \leq c$, and $\mathbf{U} \in \mathbb{R}^{n \times \rho}$ be the left singular vectors of \mathbf{D} . Let $\mathbf{Q} \in \mathbb{R}^{n \times s}$ be any orthonormal bases of $\mathbf{U}(\mathbf{U}^T \mathbf{K}^{1/2})_s$. It holds that*

$$\begin{aligned}
 (\mathbf{C} \mathbf{W}^\dagger \mathbf{C}^T)_s &= \mathbf{K}^{1/2} \mathbf{Q} \mathbf{Q}^T \mathbf{K}^{1/2}, \\
 \|\mathbf{K} - (\mathbf{C} \mathbf{W}^\dagger \mathbf{C}^T)_s\|_* &= \min_{\text{rank}(\mathbf{Z}) \leq s} \|\mathbf{K}^{1/2} - (\mathbf{K}^{1/2} \mathbf{P}) \mathbf{Z}\|_F^2.
 \end{aligned}$$

Proof Let $\mathbf{D} = \mathbf{K}^{1/2} \mathbf{P} = \mathbf{U} \Sigma \mathbf{V}^T \in \mathbb{R}^{n \times c}$ be the SVD of \mathbf{D} . It holds that $\mathbf{C} = \mathbf{K} \mathbf{P} = \mathbf{K}^{1/2} \mathbf{D}$ and $\mathbf{W} = \mathbf{P}^T \mathbf{K} \mathbf{P} = \mathbf{D}^T \mathbf{D}$. It follows the SVD of \mathbf{D} that

$$\mathbf{C}(\mathbf{W}^\dagger)^{1/2} = \mathbf{K}^{1/2} \mathbf{D}((\mathbf{D}^T \mathbf{D})^\dagger)^{1/2} = \mathbf{K}^{1/2} (\mathbf{U} \Sigma \mathbf{V}^T) (\mathbf{V} \Sigma^{-2} \mathbf{V}^T)^{1/2} = \mathbf{K}^{1/2} \mathbf{U} \mathbf{V}^T.$$

We need the following lemma to prove Lemma 8. The following lemma is not hard to prove.

Lemma 9 *Let the integers n, c, ρ, s satisfy $n \geq c \geq \rho \geq s$. Let $\mathbf{A} \in \mathbb{R}^{n \times \rho}$ be any matrix with rank at least s and $\mathbf{V} \in \mathbb{R}^{c \times \rho}$ has orthonormal columns. Then $(\mathbf{A} \mathbf{V}^T)_s = \mathbf{A}_s \mathbf{V}^T$.*

If $\text{rank}(\mathbf{K}^{1/2}\mathbf{U}) < s$, then obviously $\mathbf{K}^{1/2}\mathbf{U}\mathbf{V}^T = (\mathbf{K}^{1/2}\mathbf{U}\mathbf{V}^T)_s = (\mathbf{K}^{1/2}\mathbf{U})_s\mathbf{V}^T$. We thus assume $\text{rank}(\mathbf{K}^{1/2}\mathbf{U}\mathbf{V}^T) \geq s$ and can apply the above lemma to show $(\mathbf{K}^{1/2}\mathbf{U}\mathbf{V}^T)_s = (\mathbf{K}^{1/2}\mathbf{U})_s\mathbf{V}^T$. In either case, we have

$$(\mathbf{K}^{1/2}\mathbf{U}\mathbf{V}^T)_s = (\mathbf{K}^{1/2}\mathbf{U})_s\mathbf{V}^T.$$

It follows from (9) that $(\mathbf{C}(\mathbf{W}^\dagger)^{1/2})_s = (\mathbf{K}^{1/2}\mathbf{U})_s\mathbf{V}^T$. Thus

$$(\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s = (\mathbf{K}^{1/2}\mathbf{U})_s(\mathbf{U}^T\mathbf{K}^{1/2})_s = (\mathbf{K}^{1/2}\mathbf{U})_s\mathbf{U}^T\mathbf{U}(\mathbf{U}^T\mathbf{K}^{1/2})_s$$

Let $\mathbf{Q} \in \mathbb{R}^{n \times s}$ be the orthonormal bases of the rank s matrix $\mathbf{U}(\mathbf{U}^T\mathbf{K}^{1/2})_s$. Lemma 6 ensures that

$$\mathbf{Q}\mathbf{Q}^T\mathbf{K}^{1/2} = \mathbf{U}(\mathbf{U}^T\mathbf{K}^{1/2})_s.$$

It follows that

$$(\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s = \mathbf{K}^{1/2}\mathbf{Q}\mathbf{Q}^T\mathbf{K}^{1/2}, \quad (9)$$

by which the former conclusion of the lemma follows.

It is well known that $\|\mathbf{Y}\|_* = \text{tr}(\mathbf{Y})$ for any SPSD matrix \mathbf{Y} and that $\text{tr}(\mathbf{A}^T\mathbf{A}) = \|\mathbf{A}\|_F^2$ for any matrix \mathbf{A} . It follows from (9) that

$$\begin{aligned} \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &= \|\mathbf{K}^{1/2}(\mathbf{I}_n - \mathbf{Q}\mathbf{Q}^T)\mathbf{K}^{1/2}\|_* \\ &= \text{tr}\left(\mathbf{K}^{1/2}(\mathbf{I}_n - \mathbf{Q}\mathbf{Q}^T)(\mathbf{I}_n - \mathbf{Q}\mathbf{Q}^T)\mathbf{K}^{1/2}\right) \\ &= \|(\mathbf{I}_n - \mathbf{Q}\mathbf{Q}^T)\mathbf{K}^{1/2}\|_F^2 = \|\mathbf{K}^{1/2} - \mathbf{U}(\mathbf{U}^T\mathbf{K}^{1/2})_s\|_F^2; \end{aligned}$$

here the second equality holds because $\mathbf{K}^{1/2}(\mathbf{I}_n - \mathbf{Q}\mathbf{Q}^T)\mathbf{K}^{1/2}$ is SPSD and $(\mathbf{I}_n - \mathbf{Q}\mathbf{Q}^T)$ is orthogonal projection matrix; the last equality follows from Lemma 6. It follows from Lemma 5 that

$$\begin{aligned} \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &= \|\mathbf{K}^{1/2} - \mathbf{U}(\mathbf{U}^T\mathbf{K}^{1/2})_s\|_F^2 \\ &= \min_{\text{rank}(\mathbf{X}) \leq s} \|\mathbf{K}^{1/2} - \mathbf{U}\mathbf{X}\|_F^2 = \min_{\text{rank}(\mathbf{Z}) \leq s} \|\mathbf{K}^{1/2} - \mathbf{D}\mathbf{Z}\|_F^2, \end{aligned}$$

by which the latter conclusion of the lemma follows. ■

Lemma 10 formally defines the key *approximate matrix multiplication properties* (subspace embedding property and approximate orthogonality property) which—when sketching dimensions are chosen appropriately (see Table 7)—are shared by the uniform sampling, leverage score sampling, Gaussian projection, SRHT, and CountSketch sketching methods. Establishing approximate matrix multiplication bounds is the key step in proving many approximate regression and low-rank approximation results (Mahoney, 2011), and this lemma will be used crucially in the proof of Theorem 1.

Lemma 10 (Key Approximate Matrix Multiplication Properties for a Sketch) *Let $\eta, \epsilon, \delta_1, \delta_2 \in (0, 1)$ be fixed parameters. Fix $\mathbf{Y} \in \mathbb{R}^{n \times d}$ and let $\mathbf{V} \in \mathbb{R}^{n \times s}$ have orthonormal*

columns. Let $\mathbf{P} \in \mathbb{R}^{n \times c}$ be one of the sketching methods listed in Table 7. When c is larger than the quantity in the middle column of the corresponding row of Table 7, the subspace embedding property

$$\|\mathbf{V}^T \mathbf{P} \mathbf{P}^T \mathbf{V} - \mathbf{I}_s\|_2 \leq \eta$$

holds with probability at least $1 - \delta_1$. When c is larger than the quantity in the right-hand column of the corresponding row of Table 7, the matrix multiplication property

$$\|\mathbf{V}^T \mathbf{P} \mathbf{P}^T \mathbf{Y} - \mathbf{V}^T \mathbf{Y}\|_F^2 \leq \epsilon \|\mathbf{Y}\|_F^2$$

holds with probability at least $1 - \delta_2$.

Table 7: Sufficient sketching dimensions for the sketching properties of Lemma 10 to hold. The leverage score sampling referred to here is with respect to the leverage scores of the matrix \mathbf{V} described in Lemma 10. Similarly, μ is the row coherence of \mathbf{V} .

Sketching	Subspace Embedding	Matrix Multiplication
Leverage Score Sampling	$\frac{s}{\eta^2} \log \frac{s}{\delta_1}$	$\frac{s}{\epsilon \delta_2}$
Uniform Sampling	$\frac{\mu s}{\eta^2} \log \frac{s}{\delta_1}$	$\frac{\mu s}{\epsilon \delta_2}$
SRHT	$\frac{s + \log n}{\eta^2} \log \frac{s}{\delta_1}$	$\frac{s + \log n}{\epsilon \delta_2}$
Gaussian Projection	$\frac{s + \log(1/\delta_1)}{\eta^2}$	$\frac{s}{\epsilon \delta_2}$
CountSketch	$\frac{s}{\delta_1 \eta^2}$	$\frac{s}{\epsilon \delta_2}$

Proof This lemma is a reproduction of (Wang et al., 2016b, Lemma 2), and is a collation of heterogenously stated results from the literature. The random sampling estimates were established by Drineas et al. (2008); Wang et al. (2016a); Woodruff (2014). The Gaussian projection was firstly analyzed by Johnson and Lindenstrauss (1984); see Woodruff (2014) for a proof of the stated sufficient sketching dimensions. The SRHT estimates are from Drineas et al. (2011); Lu et al. (2013); Tropp (2011). The CountSketch estimates are from Meng and Mahoney (2013); Nelson and Nguyen (2013). ■

A.3. Completing the Proof of Theorem 1

Proof Let $\mathbf{K}_s = \mathbf{V}_s \mathbf{\Lambda}_s \mathbf{V}_s^T \in \mathbb{R}^{n \times n}$ be the truncated SVD of \mathbf{K} . It follows from Lemma 10 and the sketch sizes defined in Table 2 that

$$\begin{aligned} \|\mathbf{V}_s^T \mathbf{P} \mathbf{P}^T \mathbf{V}_s - \mathbf{I}_s\|_2 &\leq \eta, \\ \|(\mathbf{K}^{1/2} - \mathbf{K}_s^{1/2}) \mathbf{P} \mathbf{P}^T \mathbf{V}_s\|_F^2 &\leq \epsilon \|\mathbf{K}^{1/2} - \mathbf{K}_s^{1/2}\|_F^2 \end{aligned}$$

hold simultaneously with probability at least 0.9.

It follows from Lemma 8 that

$$\begin{aligned}
 \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &= \min_{\text{rank}(\mathbf{Z}) \leq s} \|\mathbf{K}^{1/2} - (\mathbf{K}^{1/2}\mathbf{P})\mathbf{Z}\|_F^2 \\
 &\leq \|\mathbf{K}^{1/2} - \mathbf{K}_s^{1/2}\|_F^2 + \|(\mathbf{K}^{1/2} - \mathbf{K}_s^{1/2})\mathbf{P}(\mathbf{V}_s^T\mathbf{P})^\dagger\|_F^2 \\
 &= \|\mathbf{K} - \mathbf{K}_s\|_* + \|(\mathbf{K}^{1/2} - \mathbf{K}_s^{1/2})\mathbf{P}\mathbf{P}^T\mathbf{V}_s(\mathbf{V}_s^T\mathbf{P}\mathbf{P}^T\mathbf{V}_s)^\dagger\|_F^2 \\
 &\leq \|\mathbf{K} - \mathbf{K}_s\|_* + \sigma_s^{-2}(\mathbf{V}_s^T\mathbf{P}\mathbf{P}^T\mathbf{V}_s) \|(\mathbf{K}^{1/2} - \mathbf{K}_s^{1/2})\mathbf{P}\mathbf{P}^T\mathbf{V}_s\|_F^2,
 \end{aligned}$$

where the former inequality follows from Lemma 7 (with $t = 1$). It follows that

$$\|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* \leq \|\mathbf{K} - \mathbf{K}_s\|_* + \frac{\epsilon}{(1-\eta)^2} \|\mathbf{K} - \mathbf{K}_s\|_*.$$

We let η be a constant and obtain the former claim of Theorem 1. The latter claim of Theorem 1 follows from Lemma 8. \blacksquare

Appendix B. Proof of Theorem 2

In this section, we will provide a proof of Theorem 2, our main result for kernel k -means approximation. We will start, in Sections B.1, B.2, and B.3, by presenting several technical tools of independent interest (Lemmas 11, 12, and 13, respectively). Then, in Section B.4, we will use these tools to prove Theorem 2.

The proof of Theorem 2 will proceed by combining Theorem 1 with Lemmas 11, 12, and 13. For convenience, the structure of proof is given in Figure 9.




Figure 9: The structure of the proof of Theorem 2.

We use the notation of *orthogonal projection matrix* throughout this section. A matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ is an orthogonal projection matrix if and only if there exists a matrix \mathbf{Q} with orthonormal columns such that $\mathbf{M} = \mathbf{Q}\mathbf{Q}^T$. It follows that $\mathbf{M}^T = \mathbf{M}$, $\mathbf{M}\mathbf{M} = \mathbf{M}$, $\mathbf{M} \preceq \mathbf{I}_n$ and that $\mathbf{I}_n - \mathbf{M}$ is also an orthogonal projection matrix.

B.1. Projection-Cost Preservation via Low-Rank Approximation

We first define projection-cost preservation, then show that certain low-rank approximations are projection-cost preserving. Projection-cost preservation was named and systematized

by Cohen et al. (2015), but the idea has been used in earlier works (Boutsidis et al., 2010, 2015; Feldman et al., 2013; Liang et al., 2014).

Definition 2 (Rank k Projection-Cost Preservation) Fix a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and error parameters $\epsilon_1, \epsilon_2 > 0$. The matrix $\mathbf{B} \in \mathbb{R}^{n \times d}$ has the rank k projection-cost preserving property with respect to \mathbf{A} if for any rank k orthogonal projection matrices $\mathbf{\Pi} \in \mathbb{R}^{n \times n}$,

$$(1 - \epsilon_1) \|\mathbf{A} - \mathbf{\Pi A}\|_F^2 \leq \|\mathbf{B} - \mathbf{\Pi B}\|_F^2 + \alpha \leq (1 + \epsilon_2) \|\mathbf{A} - \mathbf{\Pi A}\|_F^2,$$

for some fixed non-negative α that can depend on \mathbf{A} and \mathbf{B} but is independent of $\mathbf{\Pi}$.

The next lemma provides a way to construct a projection-cost preserving sketch of \mathbf{A} . It states that if a rank- s matrix $\tilde{\mathbf{A}}_s$ in the span of \mathbf{A} is constructed so that $\tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T$ approximates $\mathbf{A A}^T$ sufficiently well in terms of the trace norm, then $\tilde{\mathbf{A}}_s$ has the projection-cost preserving property. Observe that this lemma explicitly relates projection-cost preserving sketches to approximate matrix multiplication (Drineas et al., 2006; Mahoney, 2011), where the latter is measured with respect to the trace norm (Gittens and Mahoney, 2016). We provide a proof of Lemma 11 in Appendix D.

Lemma 11 (Projection-Cost Preservation via Low-Rank Approximation) Fix an error parameter $\epsilon \in (0, 1)$. Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and choose a rank- s matrix $\tilde{\mathbf{A}}_s$ that satisfies the following conditions:

- (i) there is an orthogonal projection matrix \mathbf{M} such that $\tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T = \mathbf{A M A}^T$, and
- (ii) $\|\mathbf{A A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T\|_* \leq (1 + \epsilon) \|\mathbf{A A}^T - \mathbf{A}_s \mathbf{A}_s^T\|_*$.

Then there exists a fixed $\alpha \geq 0$ such that, for any rank k projection matrices $\mathbf{\Pi} \in \mathbb{R}^{n \times n}$,

$$\|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\|_F^2 \leq \|(\mathbf{I}_n - \mathbf{\Pi})\tilde{\mathbf{A}}_s\|_F^2 + \alpha \leq (1 + \epsilon + \frac{k}{s}) \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\|_F^2.$$

B.2. Linear k -Means Clustering and Projection-Cost Preservation

Linear k -means clustering is formally defined by the optimization problem (1). To relate the linear k -means clustering problem to the projection-cost preservation property, we use an equivalent formulation of (1) as an optimization over the set of rank- k projection matrices of the form $\mathbf{X X}^T$, where \mathbf{X} is a *cluster indicator matrix*. This approach was adopted by Boutsidis et al. (2009, 2010, 2015); Cohen et al. (2015); Ding et al. (2005). Following them, we define the cluster indicator matrix in the following and give an example in Figure 10.

Definition 3 (Cluster Indicator Matrices) Let n and k be given, and let $\mathcal{J}_1, \dots, \mathcal{J}_k$ be a k -partition of the set $[n]$. The cluster indicator matrix corresponding to $\mathcal{J}_1, \dots, \mathcal{J}_k$ is the $n \times k$ matrix \mathbf{X} defined by

$$x_{ij} = \begin{cases} \frac{1}{\sqrt{|\mathcal{J}_j|}} & \text{if } i \in \mathcal{J}_j; \\ 0 & \text{otherwise.} \end{cases}$$

We take $\mathcal{X}_{n,k}$ to be the collection of cluster indicator matrices corresponding to all the k -partitions of $[n]$.

$$\mathbf{X}^T = \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|}
 \hline
 \frac{1}{\sqrt{4}} & 0 & \frac{1}{\sqrt{4}} & 0 & \frac{1}{\sqrt{4}} & 0 & \frac{1}{\sqrt{4}} & 0 & 0 & 0 & 0 \\
 \hline
 0 & 0 & 0 & \frac{1}{\sqrt{3}} & 0 & 0 & 0 & 0 & \frac{1}{\sqrt{3}} & 0 & \frac{1}{\sqrt{3}} \\
 \hline
 0 & \frac{1}{\sqrt{5}} & 0 & 0 & 0 & \frac{1}{\sqrt{5}} & 0 & \frac{1}{\sqrt{5}} & 0 & \frac{1}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\
 \hline
 \end{array} \quad \begin{array}{l}
 \mathcal{J}_1 = \{1, 3, 5, 7\} \\
 \mathcal{J}_2 = \{4, 9, 11\} \\
 \mathcal{J}_3 = \{2, 6, 8, 10, 12\}
 \end{array}$$

Figure 10: An example of a cluster indicator matrix $\mathbf{X} \in \mathcal{X}_{n,k}$. In this example, $n = 12$, $k = 3$, and \mathbf{X} corresponds to the indicated 3-partition $\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3$.

A cluster indicator matrix \mathbf{X} has exactly one non-zero entry in each row and has orthonormal columns. That is, because of the normalization of the non-zero entries, a cluster indicator matrix \mathbf{X} is an orthogonal matrix, but one with the additional constraint that there is only one non-zero entry in each row. (Indeed, the insight of Boutsidis et al. (2009, 2010, 2015) was that the rank-constrained optimization over all orthonormal matrices, which provides the usual PCA/SVD-based low-rank approximation, is a relaxation of the optimization over this smaller set of orthonormal matrices.) Among other things, it follows that $\mathbf{X}\mathbf{X}^T$ is a rank- k projection matrix.

Given this, assume $i \in \mathcal{J}_j$; it can be verified that the i -th row of $\mathbf{X}\mathbf{X}^T\mathbf{A} \in \mathbb{R}^{n \times d}$ is the centroid of \mathcal{J}_j :

$$(\mathbf{X}\mathbf{X}^T\mathbf{A})_{i:} = \mathbf{x}_{i:}\mathbf{X}^T\mathbf{A} = \frac{1}{|\mathcal{J}_j|} \sum_{l \in \mathcal{J}_j} \mathbf{a}_{l:}.$$

Consider the objective function of (1). When \mathbf{X} is the cluster indicator matrix corresponding to the input k -partition, this objective can be rewritten in terms of $\mathbf{X}\mathbf{X}^T$ as

$$\begin{aligned}
 \sum_{j=1}^k \sum_{i \in \mathcal{J}_j} \left\| \mathbf{a}_{i:} - \frac{1}{|\mathcal{J}_j|} \sum_{l \in \mathcal{J}_j} \mathbf{a}_{l:} \right\|_2^2 &= \sum_{j=1}^k \sum_{i \in \mathcal{J}_j} \left\| \mathbf{a}_{i:} - (\mathbf{X}\mathbf{X}^T\mathbf{A})_{i:} \right\|_2^2 \\
 &= \sum_{i=1}^n \left\| \mathbf{a}_{i:} - (\mathbf{X}\mathbf{X}^T\mathbf{A})_{i:} \right\|_2^2 = \|\mathbf{A} - \mathbf{X}\mathbf{X}^T\mathbf{A}\|_F^2.
 \end{aligned}$$

Therefore, the linear k -means clustering problem (1) is equivalent to the optimization problem

$$\operatorname{argmin}_{\mathbf{X} \in \mathcal{X}_{n,k}} \frac{1}{n} \|\mathbf{A} - \mathbf{X}\mathbf{X}^T\mathbf{A}\|_F^2. \quad (10)$$

In the remainder of this section, we use this equivalence between k -partitionings of $[n]$ and cluster indicator matrices without comment; e.g., we refer to the outputs of linear k -means clustering algorithms as cluster indicator matrices. Because the linear k -means clustering problem (10) is NP-hard, in practice it is solved by variants of Lloyd's algorithm.

Definition 4 defines an appropriate notion of approximation for linear k -means clustering algorithms. Definition 4 is an equivalent statement of Definition 1.

Definition 4 (γ -Approximate Algorithms) We call an algorithm \mathcal{A}_γ ($\gamma \geq 1$) a γ -approximate linear k -means algorithm if it produces a cluster indicator matrix $\tilde{\mathbf{X}}$ such that

$$\|\mathbf{A} - \tilde{\mathbf{X}}\tilde{\mathbf{X}}^T\mathbf{A}\|_F^2 \leq \gamma \cdot \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\mathbf{A} - \mathbf{X}\mathbf{X}^T\mathbf{A}\|_F^2$$

for any conformal matrix \mathbf{A} .

Given this, we now state Lemma 12, which relates projection-cost preservation to linear k -means clustering. Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be the input matrix and $\mathbf{B} \in \mathbb{R}^{n \times s}$ be a smaller matrix ($s < d$) satisfying the projection-cost preservation property. The projection-cost preserving property ensures that the k -partition obtained by applying linear k -means clustering to the rows of \mathbf{B} , instead of \mathbf{A} , is a good k -partition for the rows of \mathbf{A} . This lemma was established by Cohen et al. (2015). To make this paper self-contained, we provide a proof of Lemma 12 in Appendix E.

Lemma 12 (Projection-Cost Preservation for Linear k -Means) Fix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and assume $\mathbf{B} \in \mathbb{R}^{n \times s}$ satisfies the rank k projection-cost preservation property in Definition 2 with probability at least $1 - \delta$. Let the error parameters $\epsilon_1, \epsilon_2 \in (0, 1)$ be as in Definition 2. If $\tilde{\mathbf{X}}_{\mathbf{B}}$ is the result of applying a γ -approximate linear k -means clustering algorithm to the rows of \mathbf{B} , then with probability at least $1 - 2\delta$,

$$\|\mathbf{A} - \tilde{\mathbf{X}}_{\mathbf{B}}\tilde{\mathbf{X}}_{\mathbf{B}}^T\mathbf{A}\|_F^2 \leq \gamma \cdot \frac{1+\epsilon_2}{1-\epsilon_1} \cdot \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\mathbf{A} - \mathbf{X}\mathbf{X}^T\mathbf{A}\|_F^2.$$

B.3. Kernel k -Means Clustering

Given input vectors $\mathbf{a}_1, \dots, \mathbf{a}_n$, the kernel k -means clustering algorithm applies linear k -means clustering to feature vectors $\phi(\mathbf{a}_1), \dots, \phi(\mathbf{a}_n)$. For convenience, we let $\phi(\mathbf{a}_1)^T, \dots, \phi(\mathbf{a}_n)^T$ constitute the rows of the matrix Φ . Lemma 13 argues that all the information in Φ relevant to the kernel k -means clustering problem is present in the kernel matrix $\mathbf{K} = \Phi\Phi^T$. Variants of this lemma are well-known (Schölkopf and Smola, 2002).

Lemma 13 (Kernel Trick for Kernel k -Means) Let Φ be a matrix with n rows and let $\mathbf{K} = \Phi\Phi^T \in \mathbb{R}^{n \times n}$. Let $\mathbf{K} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$ be the EVD of \mathbf{K} . Then for any matrix \mathbf{X} with orthonormal columns,

$$\|\Phi - \mathbf{X}\mathbf{X}^T\Phi\|_F^2 = \|(\mathbf{V}\mathbf{\Lambda}^{1/2}) - \mathbf{X}\mathbf{X}^T(\mathbf{V}\mathbf{\Lambda}^{1/2})\|_F^2.$$

Proof Because \mathbf{X} has orthonormal columns, $\mathbf{X}\mathbf{X}^T$ is an orthogonal projection matrix, and $(\mathbf{I}_n - \mathbf{X}\mathbf{X}^T)$ is the orthogonal projection onto the complementary space. It follows that, as claimed,

$$\begin{aligned} \|\Phi - \mathbf{X}\mathbf{X}^T\Phi\|_F^2 &= \text{tr}\left(\Phi^T(\mathbf{I}_n - \mathbf{X}\mathbf{X}^T)\Phi\right) = \text{tr}\left((\mathbf{I}_n - \mathbf{X}\mathbf{X}^T)\mathbf{K}\right) \\ &= \text{tr}\left((\mathbf{I}_n - \mathbf{X}\mathbf{X}^T)\mathbf{V}\mathbf{\Lambda}^{1/2}\mathbf{\Lambda}^{1/2}\mathbf{V}(\mathbf{I}_n - \mathbf{X}\mathbf{X}^T)\right) = \|(\mathbf{V}\mathbf{\Lambda}^{1/2}) - \mathbf{X}\mathbf{X}^T(\mathbf{V}\mathbf{\Lambda}^{1/2})\|_F^2. \end{aligned}$$

The equalities are justified by, respectively, the definition of the squared Frobenius norm and the idempotence of projections; the cyclicity of the trace and the fact that $\mathbf{K} = \Phi\Phi^T$; the cyclicity of the trace and the EVD of \mathbf{K} ; and the definition of the squared Frobenius norm. ■

B.4. Completing the Proof of Theorem 2

To complete the proof of Theorem 2, we first establish Lemma 14, and then use it to prove Theorem 15, which is equivalent to Theorem 2. Let Φ be a matrix with n rows, and $\mathbf{K} = \Phi\Phi^T \in \mathbb{R}^{n \times n}$ be the corresponding kernel matrix.

Lemma 14 *Assume that $\mathbf{B} \in \mathbb{R}^{n \times s}$ satisfies*

$$\text{tr}(\mathbf{K} - \mathbf{B}\mathbf{B}^T) \leq (1 + \epsilon) \|\mathbf{K} - \mathbf{K}_s\|_* \quad (11)$$

and that there is an orthogonal projection matrix \mathbf{M} such that

$$\mathbf{B}\mathbf{B}^T = \mathbf{K}^{1/2}\mathbf{M}\mathbf{K}^{1/2}. \quad (12)$$

Use the rows of \mathbf{B} as input to a γ -approximation algorithm (see Definition 4) and let $\tilde{\mathbf{X}}_{\mathbf{B}}$ be the output cluster indicator matrix (see Definition 3), then

$$\|\Phi - \tilde{\mathbf{X}}_{\mathbf{B}}\tilde{\mathbf{X}}_{\mathbf{B}}^T\Phi\|_F^2 \leq \gamma(1 + \epsilon + \frac{k}{s}) \cdot \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\Phi - \mathbf{X}\mathbf{X}^T\Phi\|_F^2.$$

Proof Because assumptions (11) and (12) are satisfied, Lemma 11 ensures that \mathbf{B} is a rank k projection-cost preserving sketch of \mathbf{K} . That is, there exists a constant $\alpha \geq 0$ independent of $\mathbf{\Pi}$ such that

$$\|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{K}^{1/2}\|_F^2 \leq \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{B}\|_F^2 + \alpha \leq (1 + \epsilon + \frac{k}{s}) \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{K}^{1/2}\|_F^2$$

holds for any rank- k orthogonal projection matrix $\mathbf{\Pi}$. It follows by Lemma 12 that $\tilde{\mathbf{X}}_{\mathbf{B}}$ gives an almost optimal linear k -means clustering over the rows of $\mathbf{K}^{1/2}$:

$$\|\mathbf{K}^{1/2} - \tilde{\mathbf{X}}_{\mathbf{B}}\tilde{\mathbf{X}}_{\mathbf{B}}^T\mathbf{K}^{1/2}\|_F^2 \leq \gamma(1 + \epsilon + \frac{k}{s}) \cdot \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\mathbf{K}^{1/2} - \mathbf{X}\mathbf{X}^T\mathbf{K}^{1/2}\|_F^2.$$

Finally, Lemma 13 states that clustering the rows of $\mathbf{K}^{1/2}$ is equivalent to clustering the rows of Φ , so we reach the desired conclusion that

$$\|\Phi - \tilde{\mathbf{X}}_{\mathbf{B}}\tilde{\mathbf{X}}_{\mathbf{B}}^T\Phi\|_F^2 \leq \gamma(1 + \epsilon + \frac{k}{s}) \cdot \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\Phi - \mathbf{X}\mathbf{X}^T\Phi\|_F^2. \quad \blacksquare$$

Finally, we state and prove Theorem 15. Theorem 15 shows that approximate kernel k -means clustering using the Nyström method exhibits a $1 + \epsilon + \frac{k}{s}$ approximation ratio. Observe that Theorem 15 is equivalent to Theorem 2, and thus establishing it also establishes Theorem 2. Recall the Nyström approximation is $\mathbf{K} \approx \mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T$, where $\mathbf{C} = \mathbf{K}\mathbf{P}$, $\mathbf{W} = \mathbf{P}^T\mathbf{K}\mathbf{P}$, and \mathbf{P} is a sketching matrix.

Theorem 15 *Choose a sketching matrix $\mathbf{P} \in \mathbb{R}^{n \times c}$ and sketch size c consistent with Table 2. Let $\mathbf{B} \in \mathbb{R}^{n \times s}$ be any matrix satisfying $\mathbf{B}\mathbf{B}^T = (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s$. Let the cluster indicator matrix $\mathbf{X}_{\mathbf{B}}$ be the output of any γ -approximate k -means clustering algorithm applied to the rows of \mathbf{B} . With probability at least 0.9,*

$$\|\Phi - \mathbf{X}_{\mathbf{B}}\mathbf{X}_{\mathbf{B}}^T\Phi\|_F^2 \leq \gamma(1 + \epsilon + \frac{k}{s}) \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\Phi - \mathbf{X}\mathbf{X}^T\Phi\|_F^2.$$

Proof Theorem 1 shows that (11) holds with probability at least 0.9. Theorem 1 also shows that $\mathbf{B}\mathbf{B}^T = \mathbf{K}^{1/2}\mathbf{Q}\mathbf{Q}^T\mathbf{K}^{1/2}$ where \mathbf{Q} is a $n \times s$ matrix with orthonormal columns, so (12) holds surely. The theorem now follows from Lemma 14. \blacksquare

Appendix C. Proof of Theorem 4

Section C.1 analyzes rank-restricted Nystöm with power method. Section C.2 completes the proof of Theorem 4. Here we re-state the algorithm of Theorem 4: first, set any $c \geq s$ and draw a Gaussian projection matrix $\mathbf{P} \in \mathbb{R}^{n \times c}$; second, run the power iteration $\mathbf{P} \leftarrow \mathbf{K}\mathbf{P}$ for t times; third, orthogonalize $\mathbf{P} \in \mathbb{R}^{n \times c}$ to obtain $\mathbf{U} \in \mathbb{R}^{n \times c}$; fourth, compute $\mathbf{C} = \mathbf{K}\mathbf{U}$ and $\mathbf{W} = \mathbf{U}^T\mathbf{K}\mathbf{U}$; last, find any $\mathbf{B} \in \mathbb{R}^{n \times s}$ satisfying $\mathbf{B}\mathbf{B}^T = (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s$.

C.1. Rank-Restricted Nyström with Power Method

Theorem 16 analyzes the rank-restricted Nyström with power method. The theorem will be used to prove Theorem 4.

Theorem 16 *Let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be an SPSD matrix, s be the target rank, $c (\geq s)$ be arbitrary, and $\epsilon \in (0, 1)$ be an approximation parameter. Let $\mathbf{C} \in \mathbb{R}^{n \times c}$ and $\mathbf{W} \in \mathbb{R}^{c \times c}$ be computed by the above algorithm with $t = \mathcal{O}(\frac{\log(n/\epsilon)}{\log(\sigma_s/\sigma_{s+1})})$ power iterations, where σ_i is the i -th singular value of \mathbf{K} . If $c = s + \mathcal{O}(\log \frac{1}{\delta})$, then*

$$\|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* \leq (1 + \epsilon) \|\mathbf{K} - \mathbf{K}_s\|_*$$

with probability at least $1 - \delta$. If $c = s$, then this bound holds with a constant probability (that depends on s).

Proof Let $\mathbf{V}_{-s} \in \mathbb{R}^{n \times (n-s)}$ be the orthogonal complement of \mathbf{V}_s . Let \mathbf{P} be an $n \times c$ standard Gaussian matrix. Then $\mathbf{V}_{-s}^T\mathbf{P}$ is $(n-s) \times c$ standard Gaussian matrix, and $\mathbf{V}_s^T\mathbf{P}$ is $s \times c$ standard Gaussian matrix.

It is well known that the largest singular value of an $(n-s) \times c$ standard Gaussian matrix is at most $\sqrt{n-s} + \sqrt{c} + \zeta$ with probability $1 - \exp(-\zeta^2/2)$. See (Vershynin, 2010).

Consider $c > s$. The least singular value of an $s \times c$ standard Gaussian matrix is at least $\sqrt{c} - \sqrt{s} - \zeta$ with probability $1 - \exp(-\zeta^2/2)$. Combining the above results, we have that

$$\frac{\sigma_1^2(\mathbf{V}_{-s}^T\mathbf{P})}{\sigma_s^2(\mathbf{V}_s^T\mathbf{P})} \leq \frac{\sqrt{n-s} + \sqrt{c} + \sqrt{2\log(1/\delta)}}{\sqrt{c} - \sqrt{s} - \sqrt{2\log(1/\delta)}}$$

holds with probability $1 - 2\delta$.

Consider $c = s$. Rudelson and Vershynin (2008); Tao and Vu (2010) showed that the least singular value of an $s \times s$ standard Gaussian matrix \mathbf{G} satisfy

$$\mathbb{P}\{\sigma_s(\mathbf{G}) \leq \frac{\delta_1}{\sqrt{s}}\} = \delta_1 + \mathcal{O}(s^{-\tau}),$$

where $\tau > 0$ is a constant. It follows that

$$\frac{\sigma_1^2(\mathbf{V}_{-s}^T\mathbf{P})}{\sigma_s^2(\mathbf{V}_s^T\mathbf{P})} \leq \frac{\sqrt{n-s} + \sqrt{s} + \sqrt{2\log(1/\delta_2)}}{\delta_1/\sqrt{s}}$$

holds with probability $1 - \delta_1 - \delta_2 - \mathcal{O}(s^{-\tau})$.

For either $c > s$ or $c = s$, it follows from Lemma 17, which will be proved subsequently, that

$$\begin{aligned} \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &\leq \|\mathbf{K} - \mathbf{K}_s\|_* + \frac{\sigma_1^2(\mathbf{V}_{-s}^T\mathbf{P})}{\sigma_s^2(\mathbf{V}_s^T\mathbf{P})} \left(\frac{\sigma_{s+1}(\mathbf{K})}{\sigma_s(\mathbf{K})}\right)^{2t} \|\mathbf{K} - \mathbf{K}_s\|_* \\ &\leq \|\mathbf{K} - \mathbf{K}_s\|_* + \mathcal{O}(ns) \left(\frac{\sigma_{s+1}(\mathbf{K})}{\sigma_s(\mathbf{K})}\right)^{2t} \|\mathbf{K} - \mathbf{K}_s\|_* \end{aligned}$$

holds with constant probability, by which the theorem follows. \blacksquare

Lemma 17 *Let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be any fixed SPSP matrix, s and t be any positive integer, and $\mathbf{V}_s \in \mathbb{R}^{n \times s}$ be the top s singular vectors of \mathbf{K} . Let $\mathbf{P} \in \mathbb{R}^{n \times c}$ satisfy that $\mathbf{P}^T\mathbf{V}_s \in \mathbb{R}^{c \times s}$ has full column rank. Let $\mathbf{U} \in \mathbb{R}^{n \times c}$ be the orthonormal bases of $\mathbf{K}^t\mathbf{P}$. Let $\mathbf{C} = \mathbf{K}\mathbf{U}$ and $\mathbf{W} = \mathbf{U}^T\mathbf{K}\mathbf{U}$. Then*

$$\|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* \leq \|\mathbf{K} - \mathbf{K}_s\|_* + \left(\frac{\sigma_{s+1}(\mathbf{K})}{\sigma_s(\mathbf{K})}\right)^{2t} \frac{\sigma_1^2(\mathbf{V}_{-s}^T\mathbf{P})}{\sigma_s^2(\mathbf{V}_s^T\mathbf{P})} \|\mathbf{K} - \mathbf{K}_s\|_*.$$

Here $\mathbf{V}_{-s} \in \mathbb{R}^{n \times (n-s)}$ is the orthogonal complement of \mathbf{V}_s .

Proof It follows from Lemma 8 that

$$\begin{aligned} \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &= \min_{\text{rank}(\mathbf{Z}) \leq s} \|\mathbf{K}^{1/2} - (\mathbf{K}^{1/2}\mathbf{U})\mathbf{Z}\|_F^2 \\ &= \min_{\text{rank}(\mathbf{Z}) \leq s} \|\mathbf{A} - (\mathbf{A}\mathbf{U})\mathbf{Z}\|_F^2 \end{aligned}$$

where we let $\mathbf{A} = \mathbf{K}^{1/2}$. By definition, \mathbf{U} contains the orthonormal bases of $\mathbf{K}^t\mathbf{P} = \mathbf{A}^{2t}\mathbf{P}$, and thus there exists an $s \times s$ non-singular matrix \mathbf{R} such that $\mathbf{A}^{2t}\mathbf{P} = \mathbf{U}\mathbf{R}$. It follows that

$$\begin{aligned} \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &= \min_{\text{rank}(\mathbf{Z}) \leq s} \|\mathbf{A} - (\mathbf{A}\mathbf{A}^{2t}\mathbf{P}\mathbf{R}^{-1})\mathbf{Z}\|_F^2 \\ &= \min_{\text{rank}(\mathbf{Y}) \leq s} \|\mathbf{A} - (\mathbf{A}^{2t+1}\mathbf{P})\mathbf{Y}\|_F^2, \end{aligned}$$

where the latter equality follows from that the column spaces of $\mathbf{A}^{2t+1}\mathbf{P}$ and $\mathbf{A}^{2t+1}\mathbf{P}\mathbf{R}^{-1}$ are the same. It follows from Lemma 7 that

$$\begin{aligned} \|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* &= \min_{\text{rank}(\mathbf{Y}) \leq s} \|\mathbf{A} - (\mathbf{A}^{2t+1}\mathbf{P})\mathbf{Y}\|_F^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_s\|_F^2 + \left(\frac{\sigma_{s+1}(\mathbf{A})}{\sigma_s(\mathbf{A})}\right)^{2t} \left\| (\mathbf{A} - \mathbf{A}_s)\mathbf{P}(\mathbf{V}_s^T\mathbf{P})^\dagger \right\|_F^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_s\|_F^2 + \left(\frac{\sigma_{s+1}(\mathbf{A})}{\sigma_s(\mathbf{A})}\right)^{2t} \|\mathbf{A} - \mathbf{A}_s\|_F^2 \|\mathbf{V}_{-s}^T\mathbf{P}\|_2^2 \left\| (\mathbf{V}_s^T\mathbf{P})^\dagger \right\|_2^2, \end{aligned}$$

where we define $\mathbf{V}_{-s} \in \mathbb{R}^{n \times (n-s)}$ as the orthogonal complement of \mathbf{V}_s . By our definition $\mathbf{A} = \mathbf{K}^{1/2}$, it follows that

$$\begin{aligned} &\|\mathbf{K} - (\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T)_s\|_* \\ &\leq \|(\mathbf{K} - \mathbf{K}_s)^{1/2}\|_F^2 + \left(\frac{\sigma_{s+1}(\mathbf{K})}{\sigma_s(\mathbf{K})}\right)^{2t} \frac{\sigma_1^2(\mathbf{V}_{-s}^T\mathbf{P})}{\sigma_s^2(\mathbf{V}_s^T\mathbf{P})} \|(\mathbf{K} - \mathbf{K}_s)^{1/2}\|_F^2 \\ &= \|\mathbf{K} - \mathbf{K}_s\|_* + \left(\frac{\sigma_{s+1}(\mathbf{K})}{\sigma_s(\mathbf{K})}\right)^{2t} \frac{\sigma_1^2(\mathbf{V}_{-s}^T\mathbf{P})}{\sigma_s^2(\mathbf{V}_s^T\mathbf{P})} \|\mathbf{K} - \mathbf{K}_s\|_*, \end{aligned}$$

by which the lemma follows. ■

C.2. Completing the Proof of Theorem 4

Finally, we state and prove Theorem 18. Observe that Theorem 18 is equivalent to Theorem 4, and thus establishing it also establishes Theorem 4.

Theorem 18 *Let c , t , and $\mathbf{B} \in \mathbb{R}^{n \times s}$ be defined in the begining of this section. Set $t = \mathcal{O}(\frac{\log(n/\epsilon)}{\log(\sigma_s/\sigma_{s+1})})$. Let the cluster indicator matrix $\mathbf{X}_{\mathbf{B}}$ be the output of any γ -approximate k -means clustering algorithm applied to the rows of \mathbf{B} . If $c = s + \mathcal{O}(\log \frac{1}{\delta})$, then*

$$\|\Phi - \mathbf{X}_{\mathbf{B}} \mathbf{X}_{\mathbf{B}}^T \Phi\|_F^2 \leq \gamma(1 + \epsilon + \frac{k}{s}) \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\Phi - \mathbf{X} \mathbf{X}^T \Phi\|_F^2$$

holds with probability at least $1 - \delta$. If $c = s$, then the above inequality holds with probability $0.9 - \mathcal{O}(s^{-\tau})$, where τ is a constant.

Proof Since $\mathbf{B} \mathbf{B}^T = (\mathbf{C} \mathbf{W}^\dagger \mathbf{C}^T)_s$ by definition, Theorem 16 ensures that (11) holds with some constant probability. Theorem 1 shows that $\mathbf{B} \mathbf{B}^T = \mathbf{K}^{1/2} \mathbf{Q} \mathbf{Q}^T \mathbf{K}^{1/2}$ where \mathbf{Q} is a $n \times s$ matrix with orthonormal columns, so (12) holds surely. The theorem now follows from Lemma 14. ■

Appendix D. Proof of Lemma 11

To prove Lemma 11, we first establish a key lemma, Lemma 19. This lemma will make use of the following two assumptions.

Assumption 1 *Assume that $\tilde{\mathbf{A}}_s$ satisfies $\text{tr}(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T) \leq (1 + \epsilon) \|\mathbf{A} \mathbf{A}^T - \mathbf{A}_s \mathbf{A}_s^T\|_*$.*

Assumption 2 *Assume there exists such an orthogonal projection matrix \mathbf{M} that $\tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T = \mathbf{M} \mathbf{A} \mathbf{A}^T$.*

Lemma 19 *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be any fixed matrix. Fix an error parameter $\epsilon \in (0, 1)$. Let the rank s matrix $\tilde{\mathbf{A}}_s$ satisfy Assumptions 1 and 2. Then for any rank k orthogonal projection matrix $\mathbf{\Pi} \in \mathbb{R}^{n \times n}$,*

$$\text{tr}(\mathbf{\Pi}(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T) \mathbf{\Pi}) \leq (\epsilon + \frac{k}{s}) \|\mathbf{A} - \mathbf{\Pi} \mathbf{A}\|_F^2.$$

Proof It holds that

$$\begin{aligned} & \text{tr}((\mathbf{I}_n - \mathbf{\Pi})(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T)) \\ &= \text{tr}((\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T) - \mathbf{\Pi}(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T)) \\ &= \text{tr}(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T) - \text{tr}(\mathbf{\Pi}(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T) \mathbf{\Pi}) \\ &\leq (1 + \epsilon) \|\mathbf{A} \mathbf{A}^T - \mathbf{A}_s \mathbf{A}_s^T\|_* - \text{tr}(\mathbf{\Pi}(\mathbf{A} \mathbf{A}^T - \tilde{\mathbf{A}}_s \tilde{\mathbf{A}}_s^T) \mathbf{\Pi}), \end{aligned}$$

where the inequality follows from Assumption 1. It can be equivalently written as

$$\begin{aligned} & \text{tr}\left(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\mathbf{\Pi}\right) \\ & \leq (1 + \epsilon)\|\mathbf{A}\mathbf{A}^T - \mathbf{A}_s\mathbf{A}_s^T\|_* - \text{tr}\left((\mathbf{I}_n - \mathbf{\Pi})(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)(\mathbf{I}_n - \mathbf{\Pi})\right). \end{aligned}$$

Let \mathbf{M} be an orthogonal projection matrix defined in Assumption 2; by this assumption, it holds that $\tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T = \mathbf{A}\mathbf{M}\mathbf{A}^T$ and $\text{rank}(\mathbf{M}\mathbf{A}) = s$. It follows that

$$\begin{aligned} & \text{tr}\left((\mathbf{I}_n - \mathbf{\Pi})(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)(\mathbf{I}_n - \mathbf{\Pi})\right) = \text{tr}\left((\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}(\mathbf{I} - \mathbf{M})\mathbf{A}^T(\mathbf{I}_n - \mathbf{\Pi})\right) \\ & = \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}(\mathbf{I} - \mathbf{M})\|_F^2 = \left\|\mathbf{A} - \underbrace{\mathbf{\Pi}\mathbf{A}}_{\text{rank } k} - \underbrace{(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\mathbf{M}}_{\text{rank } s}\right\|_F^2 \\ & \geq \min_{\text{rank}(\mathbf{Y}) \leq s+k} \|\mathbf{A} - \mathbf{Y}\|_F^2 = \|\mathbf{A} - \mathbf{A}_{s+k}\|_F^2, \end{aligned}$$

where the inequality follows from that matrix rank is subadditive function. It follows that

$$\begin{aligned} & \text{tr}\left(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\mathbf{\Pi}\right) \\ & \leq (1 + \epsilon)\|\mathbf{A}\mathbf{A}^T - \mathbf{A}_s\mathbf{A}_s^T\|_* - \|\mathbf{A} - \mathbf{A}_{s+k}\|_F^2 \\ & = (1 + \epsilon)\|\mathbf{A} - \mathbf{A}_s\|_F^2 - \|\mathbf{A} - \mathbf{A}_{s+k}\|_F^2 \\ & = \epsilon\|\mathbf{A} - \mathbf{A}_s\|_F^2 + \sum_{i=s+1}^n \sigma_i^2(\mathbf{A}) - \sum_{i=s+k+1}^n \sigma_i^2(\mathbf{A}) \\ & = \epsilon\|\mathbf{A} - \mathbf{A}_s\|_F^2 + \sum_{i=s+1}^{s+k} \sigma_i^2(\mathbf{A}) \\ & \leq \epsilon\|\mathbf{A} - \mathbf{A}_s\|_F^2 + \frac{k}{s} \sum_{i=k+1}^{s+k} \sigma_i^2(\mathbf{A}), \end{aligned}$$

where the last inequality follows by the decrease of the singular values. Finally, we obtain that

$$\text{tr}\left(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\mathbf{\Pi}\right) \leq \epsilon\|\mathbf{A} - \mathbf{A}_k\|_F^2 + \frac{k}{s}\|\mathbf{A} - \mathbf{A}_k\|_F^2 \leq \left(\epsilon + \frac{k}{s}\right)\|\mathbf{A} - \mathbf{\Pi}\mathbf{A}\|_F^2,$$

by which the lemma follows. ■

Given Lemma 19, we now provide the proof of Lemma 11.

Proof For any orthogonal projection matrix \mathbf{M} , it holds that $\mathbf{M} = \mathbf{M}\mathbf{M}$. It follows that

$$\begin{aligned} & \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\|_F^2 - \|(\mathbf{I}_n - \mathbf{\Pi})\tilde{\mathbf{A}}_s\|_F^2 = \text{tr}\left((\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\mathbf{A}^T\right) - \text{tr}\left((\mathbf{I}_n - \mathbf{\Pi})\tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T\right) \\ & = \text{tr}\left((\mathbf{I}_n - \mathbf{\Pi})(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\right) = \text{tr}\left(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T\right) - \text{tr}\left(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\right) \\ & = \alpha - \text{tr}\left(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\right), \end{aligned}$$

where the last equality follows by letting $\alpha = \text{tr}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T) \geq 0$ which is independent of $\mathbf{\Pi}$. The above equality can be equivalently written as

$$\|(\mathbf{I}_n - \mathbf{\Pi})\tilde{\mathbf{A}}_s\|_F^2 + \alpha = \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\|_F^2 + \text{tr}(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\mathbf{\Pi}). \quad (13)$$

Under Assumption 2 that $\tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T \preceq \mathbf{A}\mathbf{A}^T$, it follows from (13) that

$$\begin{aligned} & \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\|_F^2 \\ & \leq \|(\mathbf{I}_n - \mathbf{\Pi})\tilde{\mathbf{A}}_s\|_F^2 + \alpha \\ & = \|(\mathbf{I}_n - \mathbf{\Pi})\mathbf{A}\|_F^2 + \text{tr}(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\mathbf{\Pi}). \end{aligned}$$

Under Assumptions 1 and 2, we can apply Lemma 19 to bound the right-hand side:

$$\text{tr}(\mathbf{\Pi}(\mathbf{A}\mathbf{A}^T - \tilde{\mathbf{A}}_s\tilde{\mathbf{A}}_s^T)\mathbf{\Pi}) \leq (\epsilon + \frac{k}{s})\|\mathbf{A} - \mathbf{\Pi}\mathbf{A}\|_F^2,$$

by which the lemma follows. ■

Appendix E. Proof of Lemma 12

Because \mathbf{B} enjoys the projection-cost preservation property, there exists a constant $\alpha \geq 0$ such that for any rank k orthogonal projection matrices $\mathbf{\Pi}_1$ and $\mathbf{\Pi}_2$, the two inequalities

$$(1 - \epsilon_1)\|\mathbf{A} - \mathbf{\Pi}_1\mathbf{A}\|_F^2 \leq \|\mathbf{B} - \mathbf{\Pi}_1\mathbf{B}\|_F^2 + \alpha, \quad (14)$$

$$\|\mathbf{B} - \mathbf{\Pi}_2\mathbf{B}\|_F^2 + \alpha \leq (1 + \epsilon_2)\|\mathbf{A} - \mathbf{\Pi}_2\mathbf{A}\|_F^2 \quad (15)$$

hold simultaneously with probability at least $1 - 2\delta$. Let $\mathbf{\Pi}_1 = \tilde{\mathbf{X}}_B\tilde{\mathbf{X}}_B^T$ and $\mathbf{X}_A^* = \text{argmin}_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\mathbf{A} - \mathbf{X}\mathbf{X}^T\mathbf{A}\|_F^2$. It follows from (14) and the definition of γ -approximate k -means algorithm that

$$\begin{aligned} (1 - \epsilon_1)\|\mathbf{A} - \tilde{\mathbf{X}}_B\tilde{\mathbf{X}}_B^T\mathbf{A}\|_F^2 & \leq \|\mathbf{B} - \tilde{\mathbf{X}}_B\tilde{\mathbf{X}}_B^T\mathbf{B}\|_F^2 + \alpha \\ & \leq \gamma \cdot \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\mathbf{B} - \mathbf{X}\mathbf{X}^T\mathbf{B}\|_F^2 + \alpha \leq \gamma\|\mathbf{B} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{B}\|_F^2 + \alpha. \end{aligned}$$

Let $\mathbf{\Pi}_2 = \mathbf{X}_A^*\mathbf{X}_A^{*T}$. It follows from (15) that

$$\|\mathbf{B} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{B}\|_F^2 + \alpha \leq (1 + \epsilon_2)\|\mathbf{A} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{A}\|_F^2.$$

Combining the above results, we have that

$$\begin{aligned} (1 - \epsilon_1)\|\mathbf{A} - \tilde{\mathbf{X}}_B\tilde{\mathbf{X}}_B^T\mathbf{A}\|_F^2 & \leq \gamma\|\mathbf{B} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{B}\|_F^2 + \alpha \\ & \leq \gamma(\|\mathbf{B} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{B}\|_F^2 + \alpha) \leq \gamma(1 + \epsilon_2)\|\mathbf{A} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{A}\|_F^2, \end{aligned}$$

where the second inequality follows from that $\gamma \geq 1$ and $\alpha \geq 0$. It follows that

$$\|\mathbf{A} - \tilde{\mathbf{X}}_B\tilde{\mathbf{X}}_B^T\mathbf{A}\|_F^2 \leq \gamma \frac{1 + \epsilon_2}{1 - \epsilon_1} \|\mathbf{A} - \mathbf{X}_A^*\mathbf{X}_A^{*T}\mathbf{A}\|_F^2 = \gamma \frac{1 + \epsilon_2}{1 - \epsilon_1} \min_{\mathbf{X} \in \mathcal{X}_{n,k}} \|\mathbf{A} - \mathbf{X}\mathbf{X}^T\mathbf{A}\|_F^2$$

This concludes our proof.

References

- Ahmed Alaoui and Michael W. Mahoney. Fast Randomized Kernel Ridge Regression with Statistical Guarantees. In *Advances in Neural Information Processing Systems (NIPS)*. 2015.
- Daniel Aloise, Amit Deshpande, Pierre Hansen, and Preyas Popat. NP-hardness of Euclidean sum-of-squares clustering. *Machine Learning*, 75(2):245–248, 2009.
- Okan Arıkan. Compression of motion capture databases. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 890–897. ACM, 2006.
- David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, 2007.
- Pranjal Awasthi, Moses Charikar, Ravishankar Krishnaswamy, and Ali Kemal Sinop. The hardness of approximation of Euclidean k-means. *arXiv preprint arXiv:1502.03316*, 2015.
- Francis Bach. Sharp analysis of low-rank kernel matrix approximations. In *International Conference on Learning Theory (COLT)*, 2013.
- Tamara L. Berg, Alexander C. Berg, Jaety Edwards, and David A. Forsyth. Who’s in the picture. *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- Christos Boutsidis, Petros Drineas, and Michael W. Mahoney. Unsupervised Feature Selection for the k-means Clustering Problem. In *Advances in Neural Information Processing Systems (NIPS)*. 2009.
- Christos Boutsidis, Anastasios Zouzias, and Petros Drineas. Random projections for k -means clustering. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismail. Near-Optimal Column-Based Matrix Reconstruction. *SIAM Journal on Computing*, 43(2):687–717, 2014.
- Christos Boutsidis, Anastasios Zouzias, Michael W. Mahoney, and Petros Drineas. Randomized Dimensionality Reduction for k -Means Clustering. *IEEE Transactions on Information Theory*, 61(2):1045–1062, 2015.
- Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58, 2009.
- Anil Chaturvedi, J. Douglas Carroll, Paul E. Green, and John A. Rotondo. A feature-based approach to market segmentation via overlapping k-centroids clustering. *Journal of Marketing Research*, 34(3):370–377, 1997.
- Jeff Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In *Problems in Analysis, Papers dedicated to Salomon Bochner*, pages 195–199. Princeton University Press, 1969.
- Ke Chen. On coresets for k-median and k-means clustering in metric and Euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009.

- Wen-Yen Chen, Yangqiu Song, Hongjie Bai, Chih-Jen Lin, and Edward Y. Chang. Parallel Spectral Clustering in Distributed Systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):568–586, 2011.
- Radha Chitta, Rong Jin, Timothy C. Havens, and Anil K. Jain. Approximate Kernel k -means: Solution to Large Scale Kernel Clustering. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2011.
- Radha Chitta, Rong Jin, and Anil K. Jain. Efficient Kernel Clustering Using Random Fourier Features. In *IEEE International Conference on Data Mining (ICDM)*, 2012.
- Jiawei Chiu and Laurent Demanet. Sublinear Randomized Algorithms for Skeleton Decompositions. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1361–1383, 2013.
- Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In *Annual ACM Symposium on Theory of Computing (STOC)*, 2013.
- Michael B Cohen, Sam Elder, Cameron Musco, Christopher Musco, and Madalina Persu. Dimensionality reduction for k -means clustering and low rank approximation. In *Annual ACM Symposium on Theory of Computing (STOC)*, 2015.
- Michael B Cohen, Cameron Musco, and Christopher Musco. Input sparsity time low-rank approximation via ridge leverage score sampling. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2017.
- Corinna Cortes, Mehryar Mohri, and Ameet Talwalkar. On the impact of kernel approximation on learning accuracy. In *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- Sanjoy Dasgupta and Yoav Freund. Random projection trees for vector quantization. *IEEE Transactions on Information Theory*, 55(7):3229–3242, 2009.
- Inderjit S Dhillon, Yuqiang Guan, and Brian Kulis. Kernel k -means: spectral clustering and normalized cuts. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2004.
- Chris Ding, Xiaofeng He, and Horst D Simon. On the equivalence of nonnegative matrix factorization and spectral clustering. In *SIAM International Conference on Data Mining (SDM)*, 2005.
- William E Donath and Alan J Hoffman. Algorithms for partitioning graphs and computer logic based on eigenvectors of connection matrices. *IBM Technical Disclosure Bulletin*, 15(3):938–944, 1972.
- William E Donath and Alan J Hoffman. Lower bounds for the partitioning of graphs. *IBM Journal of Research and Development*, 17(5):420–425, 1973.
- Petros Drineas and Michael W. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, 2005.

- Petros Drineas and Michael W Mahoney. RandNLA: randomized numerical linear algebra. *Communications of the ACM*, 59(6):80–90, 2016.
- Petros Drineas, Ravi Kannan, and Michael W. Mahoney. Fast Monte Carlo algorithms for matrices I: Approximating matrix multiplication. *SIAM Journal on Computing*, 36(1):132–157, 2006.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM Journal on Matrix Analysis and Applications*, 30(2):844–881, September 2008.
- Petros Drineas, Michael W. Mahoney, S. Muthukrishnan, and Tamás Sarlós. Faster least squares approximation. *Numerische Mathematik*, 117(2):219–249, 2011.
- Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research*, 13:3441–3472, 2012.
- Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: constant-size coresets for k-means, PCA and projective clustering. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
- Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.
- Charless Fowlkes, Serge Belongie, Fan Chung, and Jitendra Malik. Spectral grouping using the Nyström method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):214–225, 2004.
- A. Frank and A. Asuncion. UCI machine learning repository, 2010. URL <http://archive.ics.uci.edu/ml>.
- Michael R. Garey, D. S. Johnson, and Hans S. Witsenhausen. The complexity of the generalized Lloyd-Max problem. *IEEE Transactions on Information Theory*, 28(2):255–256, 1982.
- Alex Gittens and Michael W. Mahoney. Revisiting the Nyström method for improved large-scale machine learning. *Journal of Machine Learning Research*, 17(1):3977–4041, 2016.
- Alex Gittens, Jey Kottalam, Jiyan Yang, Michael F. Ringenburt, Jatin Chhugani, Evan Racah, Mohitdeep Singh, Yushu Yao, Curt Fischer, Oliver Ruebel, Benjamin Bowen, Norman G. Lewis, Michael W. Mahoney, Venkat Krishnamurthy, and Prabhat. A Multi-Platform Evaluation of the Randomized CX Low-Rank Matrix Factorization in Spark. In *IEEE International Parallel and Distributed Processing Symposium Workshops*, 2016.
- Stephen Guattery and Gary L. Miller. On the performance of spectral graph partitioning methods. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1995.

- Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- Mark S. Handcock, Adrian E. Raftery, and Jeremy M. Tantrum. Model-based clustering for social networks. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 170(2):301–354, 2007.
- Sariel Har-Peled and Soham Mazumdar. On coresets for k-means and k-median clustering. In *Annual ACM Symposium on Theory of Computing (STOC)*, 2004.
- Robert M. Haralick and Linda G. Shapiro. Image segmentation techniques. *Computer Vision, Graphics, and Image Processing*, 29(1):100–132, 1985.
- Rong Jin, Tianbao Yang, Mehrdad Mahdavi, Y Li, and Z Zhou. Improved bounds for the Nyström method with application to kernel classification. *IEEE Transactions on Information Theory*, 59(10):6939–6949, 2013.
- William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206), 1984.
- Tapas Kanungo, David M Mount, Nathan S Netanyahu, Christine D Piatko, Ruth Silverman, and Angela Y Wu. A local search approximation algorithm for k-means clustering. In *Annual Symposium on Computational Geometry*, 2002.
- Amit Kumar, Yogish Sabharwal, and Sandeep Sen. A simple linear time $(1 + \epsilon)$ -approximation algorithm for k-means clustering in any dimensions. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2004.
- Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. Sampling methods for the Nyström method. *Journal of Machine Learning Research*, 13:981–1006, 2012.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Mu Li, Wei Bi, J.T. Kwok, and Bao-Liang Lu. Large-scale Nyström kernel matrix approximation using randomized SVD. *IEEE Transactions on Neural Networks and Learning Systems*, 26(1):152–164, Jan 2015.
- Yingyu Liang, Maria-Florina F Balcan, Vandana Kanchanapally, and David Woodruff. Improved distributed principal component analysis. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Stuart Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- Yichao Lu, Paramveer Dhillon, Dean P Foster, and Lyle Ungar. Faster ridge regression via the subsampled randomized Hadamard transform. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.

- Meena Mahajan, Prajakta Nimbhorkar, and Kasturi Varadarajan. The planar k-means problem is NP-hard. In *International Workshop on Algorithms and Computation*, pages 274–285. Springer, 2009.
- Michael W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2):123–224, 2011.
- Michael W. Mahoney and Petros Drineas. Structural properties underlying high-quality randomized numerical linear algebra algorithms. In P. Bühlmann, P. Drineas, M. Kane, and M. van de Laan, editors, *Handbook of Big Data*, pages 137–154. CRC Press, 2016.
- Jiri Matousek. On approximate geometric k-clustering. *Discrete & Computational Geometry*, 24(1):61–84, 2000.
- Xiangrui Meng and Michael W Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Annual ACM Symposium on Theory of Computing (STOC)*, 2013.
- Xiangrui Meng, Joseph Bradley, Burak Yavuz, Evan Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, DB Tsai, Manish Amde, Sean Owen, Doris Xin, Reynold Xin, Michael J. Franklin, Reza Zadeh, Mateh Zaharia, and Ameet Talwalkar. MLlib: Machine Learning in Apache Spark. *Journal of Machine Learning Research*, 17(34):1–7, 2016.
- Cameron Musco and Christopher Musco. Recursive sampling for the Nystrom method. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- John Nelson and Huy L Nguyễn. OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.
- Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: analysis and an algorithm. *Advances in Neural Information Processing Systems (NIPS)*, 2002.
- Evert J. Nyström. Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben. *Acta Mathematica*, 54(1):185–204, 1930.
- Farhad Pourkamali-Anaraki and Stephen Becker. Randomized Clustered Nyström for Large-Scale Kernel Machines. *arXiv preprint arXiv:1612.06470*, 2016.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- Mark Rudelson and Roman Vershynin. The Littlewood–Offord problem and invertibility of random matrices. *Advances in Mathematics*, 218(2):600–633, 2008.
- Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002.
- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.

- Roded Sharan, Rani Elkon, and Ron Shamir. Cluster analysis and its applications to gene expression data. In *Bioinformatics and Genome Analysis*, pages 83–108. Springer, 2002.
- Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- Si Si, Cho-Jui Hsieh, and Inderjit Dhillon. Memory efficient kernel approximation. In *International Conference on Machine Learning (ICML)*, 2014.
- Mingjun Song and Sanguthevar Rajasekaran. Fast algorithms for constant approximation k-means clustering. *Transactions on Machine Learning and Data Mining*, 3(2):67–79, 2010.
- Daniel A. Spielman and Shang-Hua Teng. Spectral partitioning works: Planar graphs and finite element meshes. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 1996.
- Alexander Strehl and Joydeep Ghosh. Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3(Dec):583–617, 2002.
- Terence Tao and Van Vu. Random matrices: The distribution of the smallest singular values. *Geometric And Functional Analysis*, 20(1):260–297, 2010.
- Joel A Tropp. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis*, 3(01n02):115–126, 2011.
- Joel A Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Fixed-rank approximation of a positive-semidefinite matrix from streaming data. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- L. H. Ungar and D. P. Foster. Clustering Methods for Collaborative Filtering. In *Proc. Recommender Systems, Papers from 1998 Workshop*, 1998. Technical Report WS-98-08.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Shusen Wang and Zhihua Zhang. Improving CUR matrix decomposition and the Nyström approximation via adaptive sampling. *Journal of Machine Learning Research*, 14:2729–2769, 2013.
- Shusen Wang, Luo Luo, and Zhihua Zhang. SPSD matrix approximation via column selection: Theories, algorithms, and extensions. *Journal of Machine Learning Research*, 17(49):1–49, 2016a.
- Shusen Wang, Zhihua Zhang, and Tong Zhang. Towards more efficient SPSD matrix approximation and CUR matrix decomposition. *Journal of Machine Learning Research*, 17(210):1–49, 2016b.

- Shusen Wang, Alex Gittens, and Michael W Mahoney. Sketched ridge regression: Optimization perspective, statistical perspective, and model averaging. In *International Conference on Machine Learning (ICML)*, 2017.
- Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- Christopher Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, 2001.
- David P Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs random Fourier features: A theoretical and empirical comparison. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. *HotCloud*, 10(10-10):95, 2010.
- Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J. Franklin, Scott Shenker, and Ion Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 2012.
- Hongyuan Zha, Xiaofeng He, Chris Ding, Ming Gu, and Horst D Simon. Spectral relaxation for k-means clustering. In *Advances in Neural Information Processing Systems (NIPS)*, 2002.
- Kai Zhang and James T. Kwok. Clustered Nyström method for large scale manifold learning and dimension reduction. *IEEE Transactions on Neural Networks*, 21(10):1576–1587, 2010.
- Zhihua Zhang. The singular value decomposition, applications and beyond. *arXiv:1510.08532*, 2015.