

Chapter 5

Approximating Spectral Clustering via Sampling: A Review



Nicolas Tremblay and Andreas Loukas

5.1 Introduction

Clustering is a cornerstone of our learning process and, thus, of our understanding of the world. Indeed, we can all distinguish between a rose and a tulip precisely because we have learned what these flowers *are*. Plato would say that we learned the Idea—or Form [120]—of both the rose and the tulip, which then enables us to recognize all instances of such flowers. A machine learner would say that we learned two *classes*: their most discriminating features (shape, size, number of petals, smell, etc.) as well as their possible intra-class variability.

Mathematically speaking, the first step on the road to classifying objects (such as flowers) is to create an abstract representation of these objects: with each object i we associate a feature vector $\mathbf{p}_i \in \mathbb{R}^d$, where the dimension d of the vector corresponds to the number of features one chooses to select for the classification task. The space \mathbb{R}^d in this context is sometimes called the *feature space*. The choice of representation will obviously have a strong impact on the subsequent classification performance. Say that in the flower example we choose to represent each flower by only $d = 3$ features: the average color of each RGB channel (level of red, green, and blue) of its petals. This choice is not fail-proof: even though the archetype of the rose is red and the archetype of the tulip is yellow, we know that some varieties of both flowers can have very similar colors and thus a classification solely based on the color will necessarily lead to confusion. In fact, there are many

N. Tremblay (✉)

CNRS, University of Grenoble Alpes, GIPSA-Lab, Saint-Martin-d'Hères, France

e-mail: nicolas.tremblay@grenoble-inp.fr

A. Loukas

Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

e-mail: andreas.loukas@epfl.ch

© Springer Nature Switzerland AG 2020

F. Ros, S. Guillaume (eds.), *Sampling Techniques for Supervised or Unsupervised Tasks*, Unsupervised and Semi-Supervised Learning,

https://doi.org/10.1007/978-3-030-29349-9_5

different ways of choosing features: from features based on the expert knowledge of a botanist to features learned by a deep learning architecture from many instances of labeled images of roses and tulips, via features obtained by hybrid methods more-or-less based on human intelligence (such as the first few components of a principal component analysis of expert-based features).

The second step on the road to classifying n objects is to choose a machine learning algorithm that groups the set of n points $\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ in k classes (k may be known in advance or determined by the algorithm itself). Choosing an appropriate algorithm depends on the context:

- **Availability of pre-labeled data.** Classifying the points \mathbf{P} in k classes may be seen as assigning a label (such as “rose” or “tulip” in our $k = 2$ example) to each of the points. If one has access to some pre-labeled data, we are in the case of *supervised learning*: a more-or-less parametrized model is first learned from the pre-labeled data and then applied to the unlabeled points that need classification. If one does not have access to any pre-labeled data, we are in the case of *unsupervised learning* where classes are typically inferred only via geometrical consideration of the distribution of points in the feature space. If one has only access to a few labeled data, we are in the in-between case of *semi-supervised learning* where the known labels are typically propagated in one form or another in the feature space.
- **Inductive vs transductive learning.** Another important characteristic of a classification algorithm is whether it can be used to classify only the set of points \mathbf{P} at hand (transductive), or if it can also be directly used to classify any never-seen data point \mathbf{p}_{n+1} (inductive).

This chapter focuses on the family of algorithms jointly referred to as *spectral clustering*. These algorithms are unsupervised and transductive: no label is known in advance and one may not naturally¹ extend the results obtained on \mathbf{P} to never-seen data points. Another particularity of spectral clustering algorithms is that the number of classes k is known in advance.

Spectral clustering algorithms have received a large attention in the last two decades due to their good performance on a wide range of different datasets, as well as their ease of implementation. In a nutshell, they combine three steps:

1. **Graph construction.** A sparse similarity graph is built between the n points.
2. **Spectral embedding.** The first k eigenvectors of a graph representative matrix (such as the Laplacian) are computed.
3. **Clustering.** k -means is performed on these spectral features, to obtain k clusters.

For background information about spectral clustering, such as several justifications of its performance, out-of-sample extensions, as well as comparisons with local methods, the interested reader is referred to the recent book chapter [144].

¹Out-of-sample extensions of spectral clustering do exist (see, for instance, Section 5.3.6 of [144]), but they require additional work.

One of the drawbacks of spectral clustering is its computational cost as n , d , and/or k become large (see Sect. 5.2.3 for a discussion on the cost). Since the turn of the century, a large number of authors have striven to reduce the computational cost while keeping the high level of classification performance. The majority of such accelerating methods are based on sampling: they reduce the dimension of a subproblem of spectral clustering, compute a low-cost solution in small dimension, and lift the result back to the original space.

The goal of this chapter is to review existing sampling methods for spectral clustering, focusing especially on their approximation guarantees. Some of the fundamental questions we are interested in are: *where is the sampling performed and what is sampled precisely? how should the reduced approximate solutions be lifted back to the original space? what is the computational gain? what is the control on performances—if it exists?* Given the breadth of the literature on the subject, we do not try to be exhaustive, but rather to illustrate the key ways that sampling can be used to provide acceleration, paying special attention on recent developments on the subject.

Chapter Organization We begin by recalling in Sect. 5.2 the prototypical spectral clustering algorithm. We also provide some intuitive and formal justification of why it works. The next three sections classify the different methods of the literature depending on where the sampling is performed with respect to the three steps of spectral clustering:

- Section 5.3 details methods that sample directly in the original feature space.
- Section 5.4 assumes that the similarity graph is given and details methods that sample nodes and/or edges to approximate the spectral embedding.
- Section 5.5 assumes that the spectral embedding is given and details methods to accelerate the k -means step.

Finally, Sect. 5.6 gives perspective on the limitations of existing works and discusses key open problems.

Notation Scalars, such as λ or d , are written with lowercase letters. Vectors, such as \mathbf{u} , \mathbf{z} , or the all-one vector $\mathbf{1}$, are denoted by lowercase bold letters. Matrices, such as \mathbf{W} or \mathbf{L} , are denoted by bold capital letters. Ensembles are denoted by serif font capital letters, such as \mathbf{C} or \mathbf{X} . The “tilde” will denote approximations, such as in $\tilde{\mathbf{z}}$ or $\tilde{\mathbf{U}}_k$. We use so-called Matlab notations to slice matrices: given a set of indices \mathbf{S} of size m and an $n \times n$ matrix \mathbf{W} , $\mathbf{W}(\mathbf{S}, :)$ $\in \mathbb{R}^{m \times n}$ is \mathbf{W} reduced to the lines indexed by \mathbf{S} , $\mathbf{W}(:, \mathbf{S}) \in \mathbb{R}^{n \times m}$ is \mathbf{W} reduced to the columns indexed by \mathbf{S} , and $\mathbf{W}(\mathbf{S}, \mathbf{S}) \in \mathbb{R}^{m \times m}$ is \mathbf{W} reduced to the lines and columns indexed by \mathbf{S} . The equation $\mathbf{U}_k = \mathbf{U}(:, :k)$ defines \mathbf{U}_k as the reduction of \mathbf{U} to its first k columns. Also, \mathbf{C}^\top is the transpose of matrix \mathbf{C} and \mathbf{C}^+ its Moore–Penrose pseudo-inverse. The operator $\mathbf{X} = \text{diag}(\mathbf{x})$ takes as an input a vector $\mathbf{x} \in \mathbb{R}^n$ and returns an $n \times n$ diagonal matrix \mathbf{X} featuring \mathbf{x} in its main diagonal, i.e., $\mathbf{X}(i, j) = \mathbf{x}(i)$ if $i = j$ and $\mathbf{X}(i, j) = 0$, otherwise. Finally, we will consider graphs in a large part of this paper. We will denote by $G = (\mathbf{V}, \mathbf{E}, \mathbf{W})$ the undirected weighted graph of $|\mathbf{V}| = n$ nodes interconnected by $|\mathbf{E}| = e$ edges. $e_{ij} \in \mathbf{E}$ is the edge connecting

nodes v_i and v_j , with weight $\mathbf{W}(i, j) \geq 0$. Matrix \mathbf{W} is the adjacency matrix of G . As G is considered undirected, \mathbf{W} is also symmetric. In general, \mathbf{W} can be any symmetric matrix with positive entries, but we usually prefer to work with sparse graphs without self-loops, in which case the matrix is also sparse and has a zero diagonal.

5.2 Spectral Clustering

The input of spectral clustering algorithms consists of (1) a set of points $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n)$ (also called featured vectors) representing n objects in a feature space of dimension d , and (2) the number of classes k in which to classify these objects. The output is a partition of the n objects in k disjoint clusters. The prototypical spectral clustering algorithm [102, 121] dates back in fact to fundamental ideas by Fiedler [46] and entails the following steps:

Algorithm 1. The prototypical spectral clustering algorithm

Input. A set of n points $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n)$ in dimension d and a number of desired clusters k .

1. Graph construction (optional)
 - (a) Compute the kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$: $\forall(i, j), \mathbf{K}(i, j) = \kappa(\|\mathbf{p}_i - \mathbf{p}_j\|_2)$.
 - (b) Compute $\mathbf{W} = s(\mathbf{K})$, a sparsified version of \mathbf{K} .
 - (c) Interpret \mathbf{W} as the adjacency matrix of a weighted undirected graph G .
2. Spectral embedding
 - (a) Compute the eigenvectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k$ associated with the k smallest eigenvalues of a graph representative matrix \mathbf{R} (usually a Laplacian) computed from \mathbf{W} .
 - (b) Set $\mathbf{U}_k = [\mathbf{u}_1 | \mathbf{u}_2 | \dots | \mathbf{u}_k] \in \mathbb{R}^{n \times k}$.
 - (c) Embed the i -th node to $\mathbf{x}_i = \frac{\mathbf{U}_k(i,:)^{\top}}{q(\|\mathbf{U}_k(i,:)\|_2)}$, with $q(\cdot)$ a normalizing function.
3. Clustering
 - (a) Use k -means on $\mathbf{x}_1, \dots, \mathbf{x}_n$ in order to identify k centroids $\mathbf{c}_1, \dots, \mathbf{c}_k$.
 - (b) Voronoi tessellation: construct one cluster per centroid \mathbf{c}_ℓ and assign each object i to the cluster of the centroid closest to \mathbf{x}_i .

Output: A partition of the n points in k clusters.

A few comments are in order:

- A common choice of kernel in step 1a is the radial basis function (RBF) kernel $\kappa(\|\mathbf{p}_i - \mathbf{p}_j\|_2) = \exp(-\|\mathbf{p}_i - \mathbf{p}_j\|_2^2/\sigma^2)$ for some user-defined σ . The sparsification s of \mathbf{K} usually entails setting the diagonal to 0 and keeping only the k largest entries of each column (i.e., set all others to 0). The obtained matrix \mathbf{K}_{sp} is not symmetric in general and a final “symmetrization” step $\mathbf{W} = \mathbf{K}_{\text{sp}} + \mathbf{K}_{\text{sp}}^\top$ is necessary to obtain a matrix \mathbf{W} interpretable as the adjacency matrix of a weighted undirected graph² $G = (\mathbf{V}, \mathbf{E}, \mathbf{W})$. This graph is called the k nearest neighbor (k -NN) similarity graph (note that the k used in this paragraph has nothing to do with the number of clusters). Other kernel functions κ and sparsification methods are possible (see Section 2 of [138] for instance).
- There are several possibilities for choosing the graph representative matrix \mathbf{R} in step 2a. We consider three main choices [138]. Let us denote by \mathbf{D} the diagonal degree matrix such that $\mathbf{D}(i, i) = \sum_j \mathbf{W}(i, j)$ is the (weighted) degree of node v_i . We define the *combinatorial* graph Laplacian matrix $\mathbf{L} = \mathbf{D} - \mathbf{W}$, the *normalized* graph Laplacian matrix $\mathbf{L}_n = \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$, and the *random walk* Laplacian $\mathbf{L}_{rw} = \mathbf{I} - \mathbf{D}^{-1}\mathbf{W}$. Other popular choices include³ the non-backtracking matrix [73], degree-corrected versions of the modularity matrix [2], the Bethe–Hessian matrix [114], or similar deformed Laplacians [34].
- The normalizing function $q(\cdot)$ used in step 2c depends on which representative matrix is chosen. In the case of the Laplacians, experimental evidence as well as some theoretical arguments [138] support using a unit norm normalization for the eigenvectors of \mathbf{L}_n (i.e., q is the identity function), and no normalization for the eigenvectors of \mathbf{L} and \mathbf{L}_{rw} (i.e., $q(\cdot) = 1$).
- Step 1 of the algorithm is “optional” in the sense that in some cases the input is not a set of points but directly a graph. For instance, it could be a graph representing a social network between n individuals, where each node is an individual and there is an edge between two nodes if they know each other. The weight on each edge can represent the strength of their relation (for instance, close to 0 if they barely know each other, and close to 1 if they are best friends). The goal is then to classify individuals based on the structure of these social connections and is usually referred to as *community detection* in this context [47]. Given the input graph, and the number k of communities to identify, one can run spectral algorithms starting directly at step 2. Readers only interested in such applications can skip Sect. 5.3, which is devoted to sampling techniques designed to accelerate step 1.

²Each node v_i of \mathbf{V} represents a point \mathbf{p}_i , an undirected edge exists between nodes v_i and v_j if and only if $\mathbf{W}(i, j) \neq 0$, and the weight of that connection is $\mathbf{W}(i, j)$.

³In some of these examples, the k largest eigenvalues (instead of the k lowest in the Laplacian cases) of the representative matrix, and especially their corresponding eigenvectors, are of interest. This is only a matter of sign of the matrix \mathbf{R} and has no impact on the general discussion.

After the spectral embedding $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ has been identified, spectral clustering uses k -means in order to find the set of k centroids $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_k)$ that best represents the data. Formally, the k -means cost function to minimize reads:

$$f(\mathbf{C}; \mathbf{X}) = \sum_{\mathbf{x} \in \mathbf{X}} \min_{\mathbf{c} \in \mathbf{C}} \|\mathbf{x} - \mathbf{c}\|_2^2. \quad (5.1)$$

We would ideally hope to identify the set of k centroids \mathbf{C}^* minimizing $f(\mathbf{C}; \mathbf{X})$. Solving exactly this problem is NP-hard [42], so one commonly resorts to approximation and heuristic solutions (see, for instance, [128] for details on different such heuristics). The most famous is the so-called Lloyd-Max heuristic algorithm:

Algorithm 2. The Lloyd-Max algorithm [87]

Input. Set of n points $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ and number of desired clusters k .

1. Start from an initial guess \mathbf{C}_{ini} of k centroids
2. Iterate until convergence:
 - (a) Assign each point \mathbf{x}_i to its closest centroid to obtain a partition of \mathbf{X} in k clusters.
 - (b) Move each centroid \mathbf{c}_ℓ to the average position of all points in cluster ℓ .

Output: A set of k centroids $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_k)$.

When the clusters are sufficiently separated and \mathbf{C}_{ini} is not too far from the optimal centroids, then the Lloyd-Max algorithm converges to the correct solution [75]. Otherwise, it typically ends up in a local minimum.

A Remark on Notation Two quantities of fundamental importance in spectral clustering are the eigenvalues λ_i and especially the eigenvectors \mathbf{u}_i of the graph Laplacian matrix. We adopt the graph theoretic convention of sorting eigenvalues in non-decreasing order: $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. Also, for reasons of brevity, we overload notation and use the same symbol for the spectrum of the three Laplacians \mathbf{L} , \mathbf{L}_n , and \mathbf{L}_{rw} . Thus, we advise the reader to rely on the context in order to discern which Laplacian gives rise to the eigenvalues and eigenvectors. Finally, the reader should keep in mind that the largest eigenvalue is always bounded by 2 for \mathbf{L}_n and \mathbf{L}_{rw} .

5.2.1 An Illustration of Spectral Clustering

The first two steps of the algorithm can be understood as a non-linear transformation from the initial feature space to another feature space (that we call spectral feature space or spectral embedding): a transformation of features \mathbf{p}_i in \mathbb{R}^d to spectral features \mathbf{x}_i in \mathbb{R}^k . The first natural question that arises is why do we run k -means on the spectral features $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ that are subject to parameter tuning and costly to compute, rather than directly run k -means on the original \mathbf{P} ? Figures 5.1 and 5.2 illustrate the answer.

In Fig. 5.1, we show the result of k -means directly on a set of artificial features \mathbf{P} known as the two half-moons dataset. In this example, the intuitive ground truth is that each half-moon corresponds to a class that we want to recover. Running k -means directly in this 2D feature space will necessarily output a linear separation between the two obtained Voronoi cells and will thus necessarily fail, as no straight line can separate the two half-moons.

Spectral clustering, via the computation of the spectral features of a similarity graph, transforms these original features \mathbf{P} in spectral features \mathbf{X} that are typically linearly separable by k -means: the two half-moons are successfully recovered! We illustrate this in Fig. 5.2. In the next section, we will examine a theoretical argument aiming to justify this phenomenon.

5.2.2 Justification of Spectral Clustering

A popular approach—and by no means the only one, see Sect. 5.2.2.3—to justify spectral clustering algorithms stems from its connection to graph partitioning. Suppose that the similarity graph $G = (\mathbf{V}, \mathbf{E}, \mathbf{W})$ has been obtained and we want to compute a partition⁴ $\mathcal{P} = \{\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_k\}$ of the nodes \mathbf{V} in k groups. Intuitively, a good clustering objective function should favor strongly connected nodes to end




Fig. 5.1 Left: the two half-moons synthetic dataset ($n = 500$, $d = 2$, $k = 2$). Right: k -means with $k = 2$ directly on \mathbf{P} is unsuccessful to separate the two half-moons

⁴By definition, a partition $\mathcal{P} = \{\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_k\}$ of the nodes \mathbf{V} is such that $\cup_{\ell=1, \dots, k} \mathbf{V}_\ell = \mathbf{V}$ and $\forall \ell \neq \ell', \mathbf{V}_\ell \cap \mathbf{V}_{\ell'} = \emptyset$.




Fig. 5.2 Illustration of the spectral clustering algorithm on the two half-moons dataset ($n = 500$, $d = 2$, $k = 2$). The graph is created with a RBF kernel and via a sparsification done with k -nearest neighbors (with $k = 5$). The spectral embedding is done with the two eigenvectors associated with the two smallest eigenvalues of the combinatorial Laplacian matrix \mathbf{L} . The embedding \mathbf{X} is here in practice in $1\mathbf{D}$ as the first eigenvector of \mathbf{L} is always constant and thus not discriminative (to confirm this, first show that \mathbf{L} is a PSD matrix and then prove that $\mathbf{L}\mathbf{c} = 0$ for any constant vector \mathbf{c}). Observe how the two clusters are now linearly separable in the spectral feature space. k -means on *these* features successfully recovers the two half-moons

up in the same subset, and nodes that are far apart in the graph to end up in different subsets. This intuition can be formalized with graph cuts.

Considering two groups V_1 and V_2 , define $w(V_1, V_2) = \sum_{i \in V_1} \sum_{j \in V_2} \mathbf{W}(i, j)$ to be the total weight of all links connecting V_1 to V_2 . Also, denote by \bar{V}_ℓ the complement of V_ℓ in V , such that $\frac{1}{2}w(V_\ell, \bar{V}_\ell)$ is the total weight one needs to cut in order to disconnect V_ℓ from the rest of the graph. Given these definitions, the simplest graph cut objective function, denoted by cut , is:

$$\text{cut}(\mathcal{P} = \{V_1, \dots, V_k\}) = \frac{1}{2} \sum_{\ell=1}^k w(V_\ell, \bar{V}_\ell). \tag{5.2}$$

The best partition according to the cut criterion is $\mathcal{P}^* = \text{argmin}_{\mathcal{P}} \text{cut}(\mathcal{P})$. For $k = 2$, solving this problem can be done exactly in $O(ne + n^2 \log(n))$ amortized time using the Stoer–Wagner algorithm [126] and approximated in nearly linear time [68]. Nevertheless, this criterion is not satisfactory as it often separates an individual node from the rest of the graph, with no attention to the balance of the sizes or volumes of the groups. In clustering, one usually wants to partition into groups that are “large enough.” There are two famous ways to balance the previous cost in the machine learning literature⁵: the *ratio cut* [143] and *normalized cut* [121] cost functions, respectively defined as:

⁵The reader should note that in the graph theory literature, the measure of conductance is preferred over ncut . Conductance is $\max_{\ell} w(V_\ell, \bar{V}_\ell)/w(V_\ell)$. The two measures are equivalent when $k = 2$.

$$\text{rcut}(\mathcal{P}) = \frac{1}{2} \sum_{\ell=1}^k \frac{w(\mathbf{V}_\ell, \bar{\mathbf{V}}_\ell)}{|\mathbf{V}_\ell|} \quad \text{and} \quad \text{ncut}(\mathcal{P}) = \frac{1}{2} \sum_{\ell=1}^k \frac{w(\mathbf{V}_\ell, \bar{\mathbf{V}}_\ell)}{\text{vol}(\mathbf{V}_\ell)}, \quad (5.3)$$

where $|\mathbf{V}_\ell|$ is the number of nodes in \mathbf{V}_ℓ and $\text{vol}(\mathbf{V}_\ell) = \sum_{i \in \mathbf{V}_\ell} \sum_{j \in \mathbf{V}} \mathbf{W}(i, j)$ is the so-called volume of \mathbf{V}_ℓ . The difference between them is that ncut favors clusters of large volume, whereas rcut only considers cluster size—though for a d -regular graph with unit weights the two measures match (up to multiplication by $1/d$). Unfortunately, it is hard to minimize these cost functions directly: minimizing these two balanced costs is NP-hard [121, 139] and one needs to search over the space of all possible partitions which is of exponential size.

A Continuous Relaxation Spectral clustering may be interpreted as a continuous relaxation of the above minimization problems. Without loss of generality, in the following we concentrate on relaxing the rcut minimization problem (ncut is relaxed almost identically). Given a partition $\mathcal{P} = (\mathbf{V}_1, \dots, \mathbf{V}_k)$, let us define

$$\mathbf{C} = \left(\frac{\mathbf{z}_1}{\sqrt{|\mathbf{V}_1|}} \mid \dots \mid \frac{\mathbf{z}_k}{\sqrt{|\mathbf{V}_k|}} \right) \in \mathbb{R}^{n \times k}, \quad (5.4)$$

where $\mathbf{z}_\ell \in \mathbb{R}^n$ is the indicator vector of \mathbf{V}_ℓ :

$$\mathbf{z}_\ell(i) = \begin{cases} 1 & \text{if node } i \in \mathbf{V}_\ell, \\ 0 & \text{otherwise.} \end{cases} \quad (5.5)$$

It will prove useful in the following to remark that, independently of how the partitions are chosen, we always have that $\mathbf{C}^\top \mathbf{C} = \mathbf{I}$, the identity matrix in dimension k . With this in place, the problem of minimizing rcut can be rewritten as (see discussion in [138]):

$$\min_{\mathbf{C} \in \mathbb{R}^{n \times k}} \text{tr}(\mathbf{C}^\top \mathbf{L} \mathbf{C}) \quad \text{s.t.} \quad \mathbf{C}^\top \mathbf{C} = \mathbf{I} \quad \text{and} \quad \mathbf{C} \text{ as in (5.4)} \quad (5.6)$$

To understand why this equivalence holds, one should simply note that

$$\begin{aligned} \text{tr}(\mathbf{C}^\top \mathbf{L} \mathbf{C}) &= \sum_{\ell=1}^k \frac{1}{|\mathbf{V}_\ell|} \mathbf{z}_\ell^\top \mathbf{L} \mathbf{z}_\ell = \sum_{\ell=1}^k \frac{1}{|\mathbf{V}_\ell|} \sum_{i>j} \mathbf{W}(i, j) (\mathbf{z}_\ell(i) - \mathbf{z}_\ell(j))^2 \\ &= \sum_{\ell=1}^k \frac{w(\mathbf{V}_\ell, \bar{\mathbf{V}}_\ell)}{|\mathbf{V}_\ell|} = 2 \text{rcut}(\mathcal{P}). \end{aligned}$$

Solving (5.6) is obviously still NP-hard as the only thing we have achieved is to rewrite the rcut minimization problem in matrix form. Yet, in this form, it is easier to realize that one may find an approximate solution by relaxing the discreteness

constraint “ \mathbf{C} as in (5.4).” In the absence of the hard-to-deal-with constraint, the relaxed problem is not only polynomially solvable but also possesses a closed-form solution! By the Courant–Fischer–Weyl (min-max) theorem, the solution is given by the first k eigenvectors $\mathbf{U}_k = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k]$ of \mathbf{L} :

$$\mathbf{U}_k = \arg \min_{\mathbf{C} \in \mathbb{R}^{n \times k}} \text{tr}(\mathbf{C}^\top \mathbf{L} \mathbf{C}) \quad \text{subject to} \quad \mathbf{C}^\top \mathbf{C} = \mathbf{I}.$$

This relaxation is not unique to the combinatorial Laplacian. In the same spirit, the minimum `ncut` optimization problem can be formulated in terms of the normalized Laplacian matrix \mathbf{L}_n , and the relaxed problem’s solution is given by the first k eigenvectors of \mathbf{L}_n .

A difficulty still lies before us: how do we go from a real-valued \mathbf{U}_k to a partition of the nodes? The two next subsections aim to motivate the use of k -means as a rounding heuristic. The exposition starts from the simple case when there are only two clusters ($k = 2$) before considering the general case (arbitrary k).

5.2.2.1 The Case of Two Clusters: Thresholding Suffices

For simplicity, we first consider the case of two clusters. If one constructs a partitioning \mathcal{P}_t with $\mathbf{V}_1 = \{v_i : \mathbf{u}_2(i) > t\}$ and $\mathbf{V}_2 = \{v_i : \mathbf{u}_2(i) \leq t\}$ for every level set $t \in (-1, 1)$, then it is a folklore result that

$$\text{rcut}(\mathcal{P}^*) \leq \min_t \text{rcut}(\mathcal{P}_t) \leq 2\sqrt{\text{rcut}(\mathcal{P}^*) \left(d_{\max} - \frac{\lambda_2}{2}\right)}, \quad (5.7)$$

with $\mathcal{P}^* = \arg \min_{\mathcal{P}} \text{rcut}(\mathcal{P})$ being the optimal partitioning, d_{\max} is the maximum degree of any node in \mathbf{V} , and λ_2 the second smallest eigenvalue of \mathbf{L} . The upper bound is achieved by the tree-cross-path graph constructed by Guattery and Miller [57]. In an analogous manner, if $\mathcal{P}^* = \arg \min_{\mathcal{P}} \text{ncut}(\mathcal{P})$ is the optimal partitioning w.r.t. the `ncut` cost and every \mathcal{P}_t has been constructed by thresholding the second eigenvector of \mathbf{L}_n , then

$$\text{ncut}(\mathcal{P}^*) \leq \min_t \text{ncut}(\mathcal{P}_t) \leq 2\sqrt{\text{ncut}(\mathcal{P}^*)}. \quad (5.8)$$

Inequality (5.8) can be derived as a consequence of the *Cheeger inequality*, a key result of spectral graph theory [32], which for the normalized Laplacian reads:

$$\frac{\lambda_2}{2} \leq \text{ncut}(\mathcal{P}^*) \leq \min_V \frac{w(\mathbf{V}, \bar{\mathbf{V}})}{\min\{w(\mathbf{V}), w(\bar{\mathbf{V}})\}} \leq \min_t \text{ncut}(\mathcal{P}_t) \leq \sqrt{2\lambda_2}.$$

As a consequence, we have

$$\text{ncut}(\mathcal{P}^*) \leq \min_t \text{ncut}(\mathcal{P}_t) \leq \sqrt{2\lambda_2} \leq \sqrt{4\text{ncut}(\mathcal{P}^*)} = 2\sqrt{\text{ncut}(\mathcal{P}^*)},$$

as desired. The derivation of the `rcut` bound given in (5.7) follows similarly.

5.2.2.2 More Than Two Clusters: Use k -Means

As the number of clusters k increases, the brute-force approach of testing every level set becomes quickly prohibitive. But why is k -means the right way to obtain the clusters in the spectral embedding? Though a plethora of experimental evidence advocate the use of k -means, a rigorous justification is still lacking. The interested reader may refer to [83] for an example of an analysis of spectral partitioning without k -means.

More recently, Peng et al. [107] came up with a mathematical argument showing that, if G is well clusterable and we use a k -means algorithm (e.g., [76]) which guarantees that the identified solution $\tilde{\mathbf{C}}$ abides to

$$f(\tilde{\mathbf{C}}; \mathbf{X}) \leq (1 + \epsilon)f(\mathbf{C}^*; \mathbf{X}),$$

where \mathbf{C}^* is the optimal solution of the k -means problem, then the partitioning $\tilde{\mathcal{P}}$ produced by spectral clustering when using \mathbf{L}_n has `ncut` cost provably close to that of the optimal partitioning \mathcal{P}^* . In particular, it was shown that, as long as $\lambda_{k+1} \geq ck^2\text{ncut}(\mathcal{P}^*)$, then

$$\text{ncut}(\mathcal{P}^*) \leq \text{ncut}(\tilde{\mathcal{P}}) \leq \zeta \text{ncut}(\mathcal{P}^*) \left(1 + \epsilon \frac{k^3}{\lambda_{k+1}}\right),$$

for some constants $c, \zeta > 0$ that are independent of n and k (see also [71]). Note that, using the higher-order Cheeger inequality [83] $\lambda_k/2 \leq \text{ncut}(\mathcal{P}^*)$, the condition $\lambda_{k+1} \geq ck^2\text{ncut}(\mathcal{P}^*)$ implies

$$\frac{\lambda_{k+1}}{\lambda_k} \geq \frac{ck^2}{2} = \Omega(k^2).$$

Though hopefully milder than this one,⁶ such gap assumptions are very common in the analysis of spectral clustering. Simply put, the larger the gap $\lambda_{k+1} - \lambda_k$ is, the stronger the cluster structure and the easier it is to identify a good clustering.

⁶To construct an example possibly verifying such a strong gap assumption, consider k cliques of size k connected together via only $k-1$ edges, so as to form a loosely connected chain. Even though this is a straightforward clustering problem known to be easy for spectral clustering algorithms, the above theorem's assumption implies $\lambda_{k+1} = \Omega(k^2\text{ncut}(\mathcal{P}^*)) = \Omega(k)$ which, independently of n , can only be satisfied when k is a small (recall that the eigenvalues of \mathbf{L}_n are necessarily between 0 and 2).

Besides quantifying the difficulty of the clustering problem, the gap also encodes the robustness of the spectral embedding to errors induced by approximation algorithms [36]. The eigenvectors of a *perturbed* Hermitian matrix exhibit an interesting property: instead of being arbitrary, induced changes are localized w.r.t. the eigenvalue axis, following an inverse square eigenvalue-distance law [89]. More precisely, if $\tilde{\mathbf{u}}_i$ is the i -th eigenvector after perturbation, then the inner products $(\tilde{\mathbf{u}}_i^\top \mathbf{u}_j)^2$ decrease proportionally with $|\lambda_i - \lambda_j|^2$. As such, demanding that $\lambda_{k+1} - \lambda_k$ is large is often helpful in the analysis of spectral clustering algorithms in order to ensure that the majority of useful information (contained within \mathbf{U}_k) is preserved (in $\tilde{\mathbf{U}}_k$) despite approximation errors.⁷

5.2.2.3 Choice of Relaxation

The presented relaxation approach is not unique and other relaxations could be equally valid (see, for instance, [17, 24, 112]). This relaxation has nevertheless the double advantage of being theoretically simple and computationally easy to implement. Also, justification of spectral clustering algorithms does not only come from this graph cut perspective and in fact encompasses several approaches that we will not detail here: perturbation approaches or hitting time considerations [138], a polarization theorem [23], consistency derivations [84, 135], etc. Interestingly, recent studies (for instance, [18]) on the stochastic block models have shown that spectral clustering (on other matrices than the Laplacian, such as the non-backtracking matrix [73], or the Bethe–Hessian matrix [114] or other similar deformed Laplacians [34]) perform well up to the detectability threshold of the block structure.

5.2.3 Computational Complexity Considerations

What is the computational complexity of spectral clustering as a function of the number of points n , their dimension d , and the number of desired clusters k ? Let us examine the three steps involved one by one.

The first step entails the construction of a sparse similarity graph from the input points, which is dominated by the kernel computation and costs $\mathcal{O}(dn^2)$. In the second step, given the graph G consisting of n nodes and e edges,⁸ one needs to compute the spectral embedding (step 2 of Algorithm 1). Without exploiting the special structure of a graph Laplacian—other than its sparsity that is—there are two main options:

⁷Usually, one needs to ensure that $\sum_{i \leq k, j > k} (\tilde{\mathbf{u}}_i^\top \mathbf{u}_j)^2 / k$ remains bounded.

⁸With e of the order of n if the sparsification step was well conducted.

- Using power iterations, one may identify sequentially each non-trivial eigenvector \mathbf{u}_ℓ in time $\mathcal{O}(e/\delta_\ell)$, where $\delta_\ell = \lambda_\ell - \lambda_{\ell-1}$ is the ℓ -th eigenvalue gap and e is the number of edges of the graph [136]. Computing the spectral embedding therefore takes $\mathcal{O}(ke/\delta)$ with $\delta = \min_\ell \delta_\ell$. Unfortunately, there exist graphs⁹ such that $\delta = \mathcal{O}(1/n)$, bringing the overall worst-case complexity to $\mathcal{O}(kne)$.
- The Lanczos method can be used to approximate the first k eigenvectors in roughly $\mathcal{O}(ek + nk^2)$ time. This procedure is often numerically unstable resulting to a loss of orthogonality in the computed Krylov subspace basis. The most common way to circumvent this problem is by implicit restart [26], whose computational complexity is not easily derived. The number of restarts, empirically, depends heavily on the eigenvalue distribution in the vicinity of λ_k : if λ_k is in an eigenvalue bulk, the algorithm takes longer than when λ_k is isolated. We decide to write the complexity of restarted Arnoldi as $\mathcal{O}(t(ek + nk^2))$ with t modeling the number of restarts. Note that throughout this paper, t will generically refer to a number of iterations in algorithm complexities. We refer the interested reader to [13] for an in-depth discussion of Lanczos methods.

The third step entails solving the k -means problem, typically by using the Lloyd-Max algorithm to converge to a local minimum of $f(\mathbf{C}; \mathbf{X})$. Since there is no guarantee that this procedure will find a good local minimum, it is usually rerun multiple times, starting in each case from randomly selected centroids \mathbf{C}_{ini} . The computational complexity of this third step is $\mathcal{O}(tnk^2)$, where t is a bound on the number of iterations required until convergence multiplied by the number of retries (typically 10).

5.2.4 A Taxonomy of Sampling Methods for Spectral Clustering

For the remainder of the chapter, we propose to classify sampling methods aiming at accelerating one or more of these three steps according to when they sample. If they sample before step 1, they are detailed in Sect. 5.3. Methods that assume that the similarity graph is given or well-approximated and sample between steps 1 and 2 will be found in Sect. 5.4. Finally, methods that assume that the spectral embedding has been exactly computed or well-approximated and sample before the k -means step are explained in Sect. 5.5. This classification of methods, like all classification systems, bears a few flaws. For instance, Nyström methods can be applied to both the context of Sects. 5.3 and 5.4 and are thus mentioned in both. Also, we decided to include the pseudo-code of only a few chosen algorithms that we think are illustrative of the literature. This choice is of course subjective and

⁹The combinatorial Laplacian of a complete balanced binary tree on $k \geq 3$ levels and $n = 2^k - 1$ nodes has $\frac{1}{n} \leq \lambda_2 \leq \frac{2}{n}$ [56].

debatable. Notwithstanding these flaws, we hope that this taxonomy clarifies the landscape of existing methods.

5.3 Sampling in the Original Feature Space

This section is devoted to methods that ambitiously aim to reduce the dimension of the spectral clustering problem even before the graph has been formed. Indeed, the naive way of building the similarity graph (step 1 of spectral clustering algorithms) costs $\mathcal{O}(dn^2)$ and, as such, is one of the main computational bottlenecks of spectral clustering. It should be remarked that the present discussion fits into the wider realm of kernel approximation, a proper review of which cannot fit in this chapter: we will thus concentrate on methods that were in practice used for spectral clustering.

5.3.1 Nyström-Based Methods

The methods of this section aim to obtain an approximation $\tilde{\mathbf{U}}_k$ of the exact spectral embedding \mathbf{U}_k via a sampling procedure in the original feature space.

The Nyström method is a well-known algorithm for obtaining a rough low-rank approximation of a positive semi-definite (PSD) matrix \mathbf{A} . Here is a high level description of the steps entailed:

Algorithm 3. Nyström's method

Input. PSD matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, number of samples m , desired rank k

1. Let \mathbf{S} be m column indices chosen by some sampling procedure.
2. Denote by $\mathbf{B} = \mathbf{A}(\mathbf{S}, \mathbf{S}) \in \mathbb{R}^{m \times m}$ and $\mathbf{C} = \mathbf{A}(:, \mathbf{S}) \in \mathbb{R}^{n \times m}$ the sub-matrices indexed by \mathbf{S} .
3. Let $\mathbf{B} = \mathbf{Q}\mathbf{\Sigma}\mathbf{Q}^\top$ be the eigen-decomposition of \mathbf{B} with the diagonal of $\mathbf{\Sigma}$ sorted in decreasing magnitude.
4. Compute the rank- k approximation of \mathbf{B} as $\mathbf{B}_k = \mathbf{Q}_k\mathbf{\Sigma}_k\mathbf{Q}_k^\top$, where $\mathbf{Q}_k = \mathbf{Q}(:, :k) \in \mathbb{R}^{m \times k}$ and $\mathbf{\Sigma}_k = \mathbf{\Sigma}(:, :k)$.

Possible outputs:

- A low-rank approximation $\tilde{\mathbf{A}} = \mathbf{C}\mathbf{B}_k^\dagger\mathbf{C}^\top \in \mathbb{R}^{n \times n}$ of \mathbf{A}
- A rank- k approximation $\tilde{\mathbf{A}}_k = \mathbf{C}\mathbf{B}_k^\dagger\mathbf{C}^\top \in \mathbb{R}^{n \times n}$ of \mathbf{A}
- The top k eigenvectors of $\tilde{\mathbf{A}}_k$, stacked as columns in matrix $\tilde{\mathbf{V}}_k \in \mathbb{R}^{n \times k}$, obtained by orthonormalizing the columns of $\tilde{\mathbf{Q}}_k = \mathbf{C}\mathbf{Q}_k\mathbf{\Sigma}_k^{-1} \in \mathbb{R}^{n \times k}$

Various guarantees are known for the quality of $\tilde{\mathbf{A}}$ depending on the type of sampling utilized (i.e., how the indices in \mathbf{S} are selected in step 1) and the preferred notion of error (spectral $\|\cdot\|_2$ vs Frobenius $\|\cdot\|_F$ vs trace $\|\cdot\|_*$ norm) [50, 54, 77, 148]. For instance:

Theorem 5.1 (Lemma 8 for $q = 1$ in [54]) *Let $\epsilon \in (0, 1)$ and $\delta \in (0, 1)$ and suppose that \mathbf{S} contains the indices of m columns drawn i.i.d. uniformly at random (with or without replacement). Then:*

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \leq \left(1 + \frac{n}{(1 - \epsilon)m}\right) \|\mathbf{A} - \mathbf{A}_k\|_2$$

holds with probability at least $1 - 3\delta$, provided that $m \geq 2\epsilon^{-2}\mu k \log(k/\delta)$, where

$$\mu = \frac{n}{k} \max_{i=1, \dots, n} \|\mathbf{V}_k(i, \cdot)\|_2^2$$

is the coherence associated with the first k eigenvectors \mathbf{V}_k of \mathbf{A} , and \mathbf{A}_k is the best rank- k approximation of \mathbf{A} .

Guarantees independent of the coherence can be obtained for more advanced sampling methods. Perhaps the most well-known method is that of *leverage scores*, where one draws m samples independently by selecting (with replacement) the i -th column with probability $\mathbf{p}_i = \|\mathbf{V}_k(i, \cdot)\|_2^2/k$.

Theorem 5.2 (Lemma 5 for $q = 1$ in [54]) *Let $\epsilon \in (0, 1)$ and $\delta \in (0, 1)$ and suppose that \mathbf{S} contains the indices of m columns drawn i.i.d. with replacement from such a probability distribution. Then:*

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \leq \|\mathbf{A} - \mathbf{A}_k\|_2 + \epsilon^2 \|\mathbf{A} - \mathbf{A}_k\|_*$$

holds with probability at least $0.8 - 2\delta$ provided that $m \geq \mathcal{O}(\epsilon^{-2}k \log(k/\delta))$.

Computing leverage scores exactly is computationally prohibitive since it necessitates a partial SVD decomposition of \mathbf{A} , which we are trying to avoid in the first place. Nevertheless, it is possible to approximate all leverage scores with a multiplicative error guarantee in time roughly $\mathcal{O}(ek \log(e))$ if \mathbf{A} has $\mathcal{O}(e)$ non-zero entries. (see Algorithms 1–3 in [54]). Many variants of the above exist [77, 78], but to the best of our knowledge, the fastest current Nyström algorithm utilizes ridge leverage scores with a complex recursive sampling scheme and runs in time nearly linear in n [100].

Nyström for Spectral Clustering Though initially conceived for low-rank approximation, Nyström’s method can also be used to accelerate spectral clustering. The key observation is that \mathbf{U}_k , the tailing k eigenvectors of the graph representative matrix \mathbf{R} , can be interpreted as the top k eigenvectors of the PSD matrix $\mathbf{A} = \|\mathbf{R}\|_2 \mathbf{I} - \mathbf{R}$. As such, the span of the k top eigenvectors of $\tilde{\mathbf{A}}_k$ obtained by running

Algorithm 3 on \mathbf{A} is an approximation of the span of the exact spectral embedding. Different variants of this idea have been considered for the acceleration of spectral clustering [19, 48, 85, 86, 97, 141].

Following our taxonomy, we hereby focus on the case where we have at our disposal n points \mathbf{p}_i in dimension d , and the similarity graph has yet to be formed. The case where the graph is known is deferred to Sect. 5.4.

In this case, we cannot run Algorithm 3 on $\mathbf{A} = \|\mathbf{R}\|_2 \mathbf{I} - \mathbf{R}$ as the graph, and *a fortiori* its representative matrix \mathbf{R} has not yet been formed. What we *can* have access to *efficiently* is $\mathbf{B} = s(\mathbf{K}(\mathbf{S}, \mathbf{S}))$ and $\mathbf{C} = s(\mathbf{K}(:, \mathbf{S}))$, as these require only a partial computation of the kernel and cost only $\mathcal{O}(dnm)$. Note that s is a sparsification function that is applied on a subset of the kernel matrix.

The following pseudo-code exemplifies how Nyström-based techniques can be used to approximate the first k eigenvectors \mathbf{U}_k associated with the normalized Laplacian matrix (i.e., here $\mathbf{R} = \mathbf{L}_n$):

Algorithm 3b. Nyström for spectral clustering [85]

Input. The set of points \mathbf{P} , the number of desired clusters k , a sampling set \mathbf{S} of size $m \geq k$

1. Compute the sub-matrices $\mathbf{B} = s(\mathbf{K}(\mathbf{S}, \mathbf{S})) \in \mathbb{R}^{m \times m}$ and $\mathbf{C} = s(\mathbf{K}(:, \mathbf{S})) \in \mathbb{R}^{n \times m}$, where s is a sparsification function.
2. Let $\mathbf{D}_r = \text{diag}(\mathbf{B}\mathbf{1})$ be the $m \times m$ degree matrix.
3. Compute the top k eigenvalues Σ_k and eigenvectors \mathbf{Q}_k of $\mathbf{D}_r^{-1/2} \mathbf{B} \mathbf{D}_r^{-1/2}$.
4. Set $\tilde{\mathbf{Q}}_k = \mathbf{C} \mathbf{D}_r^{-1/2} \mathbf{Q}_k \Sigma_k^{-1}$.
5. Let $\mathbf{D}_l = \text{diag}(\tilde{\mathbf{Q}}_k \Sigma_k \tilde{\mathbf{Q}}_k^\top \mathbf{1})$ be the $n \times n$ degree matrix.
6. Compute $\tilde{\mathbf{U}}_k$ obtained by orthogonalizing $\mathbf{D}_l^{-1/2} \tilde{\mathbf{Q}}_k$.

Output: $\tilde{\mathbf{U}}_k$, an approximation of the spectral embedding \mathbf{U}_k .

This algorithm runs in $\mathcal{O}(nm \max(d, k))$ time, which is small when m depends mildly on the other parameters of interest. Nevertheless, the algorithm (and others like it) suffers from several issues:

- Algorithm 3b attempts to use Nyström’s method on $\mathbf{A} = 2\mathbf{I} - \mathbf{L}_n = \mathbf{I} + \mathbf{D}^{-\frac{1}{2}} s(\mathbf{K}) \mathbf{D}^{-\frac{1}{2}}$ via the exact computation of two sub-matrices of \mathbf{K} . In doing so, it makes two strong (and uncontrolled) approximations. First of all, the sparsification step (step 1 in Algorithm 3b) is applied to the sub-matrices $\mathbf{K}(\mathbf{S}, \mathbf{S})$ and $\mathbf{K}(:, \mathbf{S})$, deviating from the correct sparsification procedure that takes into account the entire kernel matrix \mathbf{K} . Second, the degree matrix \mathbf{D} is never exactly computed as knowing it exactly would entail computing exactly $s(\mathbf{K})$, which is precisely what we are trying to avoid. Existing methods thus rely on heuristic approximations of the degree in order to bypass this difficulty (see steps 2 and 5 of Algorithm 3b).

- Since we do not have direct access to the kernel matrix, we cannot utilize advanced sampling methods such as leverage scores to draw the sampling set \mathbf{S} . This is particularly problematic if (due to sparsification) matrices \mathbf{B} and \mathbf{C} are sparse, as for sparse matrices uniform sampling is known to perform poorly [97]. Techniques that rely on distances between columns do not fair much better. Landmark-based approaches commonly perform better in simple problems but suffer when the clusters are non-convex [19]. We refer the reader to the work by Mohan et al. [97] for more information on landmark-based methods. The latter work also describes an involved sampling scheme that is aimed at general (i.e., non-convex) clusters.

For the reasons highlighted above, the low-rank approximation guarantees accompanying the classical Nyström method cannot be directly used here. *A fortiori*, it is an open question how much the quality of the spectral clustering solution is affected by using the centroids obtained by running k -means on $\tilde{\mathbf{U}}_k$.

Column Sampling Akin in spirit to Nyström methods, an alternative approach to accelerating spectral clustering was inspired by column sampling low-rank approximation techniques [37, 43].

An instance of such algorithms was put forth under the name of cSPEC (column sampling spectral clustering) by Wang et al. [141]. Let $\mathbf{C} = \mathbf{U}_C \mathbf{\Sigma}_C \mathbf{V}_C^\top$ be the singular value decomposition of the $n \times m$ matrix $\mathbf{C} = s(\mathbf{K}(:, \mathbf{S}))$. Then, matrices

$$\tilde{\mathbf{\Sigma}} = \sqrt{\frac{n}{m}} \mathbf{\Sigma}_C \quad \text{and} \quad \tilde{\mathbf{U}} = \mathbf{C} \mathbf{V}_C \mathbf{\Sigma}_C^+$$

are interpreted as an approximation of the actual eigenvalues and eigenvectors of \mathbf{K} and thus \mathbf{U}_k can be substituted by the first k columns of $\tilde{\mathbf{U}}$. This algorithm runs in $\mathcal{O}(ndm + nm^2)$.

Authors in [29] propose a hybrid method, between column sampling and the representative-based methods discussed in Sect. 5.3.3, where they propose the following approximate factorization of the data matrix:

$$(\mathbf{p}_1 | \dots | \mathbf{p}_n) \simeq \mathbf{F} \mathbf{Z} \in \mathbb{R}^{d \times n}, \quad (5.9)$$

where $\mathbf{F} \in \mathbb{R}^{d \times m}$ concatenates the feature vectors of m sampled points and $\mathbf{Z} \in \mathbb{R}^{m \times n}$ represents all unsampled points as approximate linear combinations of the representatives, computed via sparse coding techniques [82].¹⁰ The SVD of $\tilde{\mathbf{D}}^{-1/2} \mathbf{Z}$, with $\tilde{\mathbf{D}}$ the row-sum of \mathbf{Z} , is then computed to obtain an approximation $\tilde{\mathbf{U}}_k$ of \mathbf{U}_k . The complexity of their algorithm is also $\mathcal{O}(ndm + nm^2)$.

¹⁰Authors in [116] have a very similar proposition as [29], adding a projection phase at the beginning to reduce the dimension d (see Sect. 5.3.4.2). Similar ideas may also be found in [137].

In these methods, the choice of the sample set \mathbf{S} is, of course, central and has been much debated. Popular options are uniformly at random or via better-tailored probability distributions, via a first k -means (with $k = m$) pass on \mathbf{P} , or via other selective sampling methods. Also, as with most extensions of Nyström’s method to spectral clustering, column sampling methods for spectral clustering do not come with end-to-end approximation guarantees on \mathbf{U}_k .

In the world of low-rank matrix approximation the situation is somewhat more advanced. Recent work in column sampling utilizes adaptive sampling with leverage scores in time $\mathcal{O}(e + npoly(k))$, or uniformly i.i.d. after preconditioning by a fast randomized Hadamard transform [41, 145]. Others have also used a correlated version called volume sampling to obtain column indices [37]. Nevertheless, this literature extends beyond the scope of this chapter and thus we invite the interested reader to consider the aforementioned references for a more in-depth perspective.

5.3.2 Random Fourier Features

Out of several sketching techniques one could *a priori* use to accelerate spectral clustering, we focus on random Fourier features (RFF) [110]: a method that samples in the Fourier space associated with the original feature space. Even though RFFs have originally been developed to approximate a kernel matrix \mathbf{K} in time linear in n instead of the quadratic time necessary for its exact computation, they can in fact be used to obtain an approximation $\tilde{\mathbf{U}}_k$ of the exact spectral embedding \mathbf{U}_k .

Let us denote by κ the RBF kernel, i.e., $\kappa(\mathbf{t}) = \exp(-\mathbf{t}^2/\sigma^2)$, whose Fourier transform is:

$$\hat{\kappa}(\boldsymbol{\omega}) = \int_{\mathbb{R}^d} \kappa(\mathbf{t}) \exp^{-i\boldsymbol{\omega}^\top \mathbf{t}} \, d\mathbf{t}. \quad (5.10)$$

The above takes real values as κ is symmetric. One may write:

$$\kappa(\mathbf{p}, \mathbf{q}) = \kappa(\mathbf{p} - \mathbf{q}) = \frac{1}{Z} \int_{\mathbb{R}^d} \hat{\kappa}(\boldsymbol{\omega}) \exp^{i\boldsymbol{\omega}^\top (\mathbf{p} - \mathbf{q})} \, d\boldsymbol{\omega}, \quad (5.11)$$

where, in order to ensure that $\kappa(\mathbf{p}, \mathbf{p}) = 1$, the normalization constant is set to $Z = \int_{\mathbb{R}^d} \hat{\kappa}(\boldsymbol{\omega}) \, d\boldsymbol{\omega}$. According to Bochner’s theorem, and due to the fact that κ is positive-definite, $\hat{\kappa}/Z$ is a valid probability density function. $\kappa(\mathbf{p}, \mathbf{q})$ may thus be interpreted as the expected value of $\exp^{i\boldsymbol{\omega}^\top (\mathbf{p} - \mathbf{q})}$ provided that $\boldsymbol{\omega}$ is drawn from $\hat{\kappa}/Z$:

$$\kappa(\mathbf{p}, \mathbf{q}) = \mathbb{E}_{\boldsymbol{\omega}} \left(\exp^{i\boldsymbol{\omega}^\top (\mathbf{p} - \mathbf{q})} \right) \quad (5.12)$$

Drawing $\boldsymbol{\omega}$ from the distribution $\hat{\kappa}/Z$ is equivalent to drawing independently each of its d entries according to the normal law of mean 0 and variance $2/\sigma^2$. Indeed: $\hat{\kappa}(\boldsymbol{\omega}) = \pi^{d/2} \sigma^d \exp(-\sigma^2 \boldsymbol{\omega}^2/4)$ and $Z = \int_{\mathbb{R}^d} \hat{\kappa}(\boldsymbol{\omega}) \, d\boldsymbol{\omega} = (2\pi)^d$, leading to

$$\frac{\hat{\kappa}(\boldsymbol{\omega})}{Z} = \left(\frac{\sigma}{2\sqrt{\pi}} \right)^d \exp^{-\sigma^2 \boldsymbol{\omega}^2 / 4}.$$

In practice, we draw independently m such vectors $\boldsymbol{\omega}$ to obtain the set of sampled frequencies $\Omega = (\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m)$. For each data point \mathbf{p}_i , and given this set of samples Ω , we define the associated random Fourier feature vector:

$$\boldsymbol{\psi}_i = \frac{1}{\sqrt{m}} [\cos(\boldsymbol{\omega}_1^\top \mathbf{p}_i) | \dots | \cos(\boldsymbol{\omega}_m^\top \mathbf{p}_i) | \sin(\boldsymbol{\omega}_1^\top \mathbf{p}_i) | \dots | \sin(\boldsymbol{\omega}_m^\top \mathbf{p}_i)]^\top \in \mathbb{R}^{2m}, \quad (5.13)$$

and call $\boldsymbol{\Psi} = (\boldsymbol{\psi}_1 | \dots | \boldsymbol{\psi}_n) \in \mathbb{R}^{2m \times n}$ the RFF matrix. Other embeddings are possible in the RFF framework, but this one was shown to be the most appropriate to the Gaussian kernel [127]. As m increases, $\boldsymbol{\psi}_i^\top \boldsymbol{\psi}_j$ concentrates around its expected value $\kappa(\mathbf{p}_i, \mathbf{p}_j)$: $\boldsymbol{\psi}_i^\top \boldsymbol{\psi}_j \simeq \kappa(\mathbf{p}_i, \mathbf{p}_j)$. Proposition 1 of [127] states the tightness of this concentration: it shows that the approximation starts to be valid with high probability for $m \geq \mathcal{O}(d \log d)$. The Gaussian kernel matrix is thus well approximated as $\mathbf{K} \simeq \boldsymbol{\Psi}^\top \boldsymbol{\Psi}$. With such a low-rank approximation $\boldsymbol{\Psi}$ of \mathbf{K} , one can estimate the degrees,¹¹ degree-normalize $\boldsymbol{\Psi}$ to obtain a low-rank approximation of the normalized Laplacian \mathbf{L}_n , and perform an SVD to directly obtain an approximation $\tilde{\mathbf{U}}_k$ of the spectral embedding \mathbf{U}_k . The total cost to obtain this approximation is $\mathcal{O}(ndm + nm^2)$. These ideas were developed in Refs. [31, 146], for instance.

As in Nyström methods however, the concentration guarantees of RFFs for \mathbf{K} do not extend to the degree-normalized case; moreover, the sparsification step 1b of spectral clustering is ignored. Note that improving over RFFs in terms of efficiency and concentration properties is the subject of recent research (see, for instance, [81]).

5.3.3 The Paradigm of Representative Points

The methods detailed here sample in the original feature space and directly obtain a control on the misclustering rate due to the sampling process. They are based on the following framework:

1. Sample m so-called representatives.
2. Run spectral clustering on the representatives.
3. Lift the solution back to the entire dataset.

¹¹An approximation of the degree d_i of node v_i is $\boldsymbol{\psi}_i^\top \bar{\boldsymbol{\psi}}$ where $\bar{\boldsymbol{\psi}} = \sum_j \boldsymbol{\psi}_j$. All degrees can thus be estimated in time $\mathcal{O}(nm^2)$.

Let us illustrate this with the example of KASP:

Algorithm 4. KASP: k -means-based approximate spectral clustering [147]

Input. A set of n points $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n)$ in dimension d , a number of desired clusters k , and a number of representatives m .

1. Perform k -means with $k = m$ on \mathbf{P} and obtain:
 - (a) the cluster centroids $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_m)$ as the m representative points.
 - (b) a correspondence table to associate each \mathbf{p}_i to its nearest representative
2. Run spectral clustering on \mathbf{Y} to get the cluster membership of each \mathbf{y}_i .
3. Lift the cluster membership to each \mathbf{p}_i by looking up the cluster membership of its representative in the correspondence table.

Output: k clusters

The complexity of KASP is bounded by¹² $\mathcal{O}(mdnt + m^3)$. For a summary of the analysis given in [147], let us consider the cluster memberships given by exact spectral clustering on \mathbf{P} as well as the memberships given by exact spectral clustering on $\tilde{\mathbf{P}} = (\mathbf{p}_1 + \epsilon_1, \dots, \mathbf{p}_n + \epsilon_n)$, where the ϵ_i are any small perturbations on the initial points. Let us denote by \mathbf{L}_n (resp. $\tilde{\mathbf{L}}_n$) the normalized Laplacian matrix of the similarity graph on \mathbf{P} (resp. $\tilde{\mathbf{P}}$). The analysis concentrates on the study of the miss-clustering rate ρ :

$$\rho = \frac{\# \text{ of points with different memberships}}{n}. \quad (5.14)$$

The main result, building upon preliminary work in [63], stems from a perturbation approach and reads:

Theorem 5.3 *Under the assumptions of Theorem 3 in [147]: $\rho \leq \mathcal{O}\left(\frac{k}{s_0^2} \|\mathbf{L}_n - \tilde{\mathbf{L}}_n\|_F\right)$, where s_0 is a value depending on the spectral gap. Also, under the assumptions of Theorem 6 in [147], one has, with high probability:*

$$\|\mathbf{L}_n - \tilde{\mathbf{L}}_n\|_F \leq \mathcal{O}\left(\sigma_\epsilon^{(2)} + \sigma_\epsilon^{(4)}\right), \quad (5.15)$$

with $\sigma_\epsilon^{(2)}$ and $\sigma_\epsilon^{(4)}$ the 2nd and 4th moments of the perturbation's norms $\|\epsilon_i\|_2$.

¹²It is in fact $\mathcal{O}(mdnt)$ for step 1, and bounded by $\mathcal{O}(dm^2 + m^2k + mk^2)$ for step 2. As $n \geq m$ and $m \geq k$, the total complexity is bounded by $\mathcal{O}(mdnt + m^3)$.

Combining both bounds, one obtains an upper bound on the misclustering rate that depends on the second and fourth moments of the perturbation's norms $\|\epsilon_i\|_2$. The “collapse” of points onto the m representative points, interpreted as a perturbation on the original points, should thus tend to minimize these two moments, leading the authors to propose *distortion-minimizing* algorithms, such as KASP. A very similar algorithm, eSPEC, is described in [141].

5.3.4 Other Methods

5.3.4.1 Approximate Nearest Neighbor Search Algorithms

The objective here is to approximate the nearest neighbor graph efficiently. Even though these methods are not necessarily based on sampling, we include them in the discussion as they are frequently used in practice.

Given the feature vectors $\mathbf{p}_1, \dots, \mathbf{p}_n \in \mathbb{R}^d$ and a query point $\mathbf{q} \in \mathbb{R}^d$, the exact nearest neighbor search (exact NNS) associated with \mathbf{P} and \mathbf{q} is $\mathbf{p}^* = \operatorname{argmin}_{\mathbf{p} \in \mathbf{P}} \operatorname{dist}(\mathbf{q}, \mathbf{p})$, where *dist* stands for any distance. Different distances are possible depending on the choice of kernel κ . We will here consider the Euclidean norm as it enters the definition of the classical RBF kernel. Computing the exact NNS costs $\mathcal{O}(dn)$. The goal of the approximate NNS field of research is to provide faster algorithms that have the following control on the error.

Definition 5.1 Point \mathbf{p}^* is an ϵ -approximate nearest neighbor of query $\mathbf{q} \in \mathbb{R}^d$, if

$$\forall \mathbf{p} \in \mathbf{P} \quad \operatorname{dist}(\mathbf{q}, \mathbf{p}^*) \leq (1 + \epsilon) \operatorname{dist}(\mathbf{q}, \mathbf{p}).$$

For $\epsilon = 0$, this reduces to exact NNS.

Extensions of this objective to the k -nearest neighbor goal are considered in the NNS literature. A k -nearest neighbor graph can then be constructed simply by running an approximate k -NNS query for each object \mathbf{p}_i . Thus, approximate NNS algorithms are interesting candidates to approximate the adjacency matrix of the nearest-neighbor affinity graph, that we need in step 1 of spectral clustering. Many algorithms exist, their respective performances depending essentially on the dimension d of the feature vectors. According to [9], randomized k - d forests as implemented in the library FLANN [98] are considered state of the art for dimension of around 100, whereas methods based on balanced box decomposition (BBD) [4, 7] are known to perform well for d roughly smaller than 100. In high dimensions, to avoid the curse of dimensionality, successful approaches are, for instance, based on hashing methods (such as locality sensitive hashing (LSH) [5], product quantization (PQ) [66]), or k - d generalized random forests [9]. Finally, proximity graph methods that sequentially improve over a first coarse approximation of the k -NN graph (or other graph structures such as navigable graphs) have received a large attention recently and are becoming state of the art in regimes where quality of approximation

primes (see, for instance, [8, 40, 51, 94]). Such tools come with various levels of guarantees and computation costs, the details of which are not in the scope of this chapter.

Experimentally, to obtain an approximate k -NN graph with a typical recall rate¹³ of 0.9, these algorithms are observed to achieve a complexity of $\mathcal{O}(dn^\alpha)$ with α close to 1 ($\alpha \simeq 1.1$ in [40], for instance).

5.3.4.2 Feature Selection and Feature Projection

Some methods work on reducing the factor d of the complexity $\mathcal{O}(dn^2)$ of the kernel computation via feature selection, i.e., the sampling of features deemed more useful for the underlying clustering task, or feature projection, i.e., the projection on usually random subspaces of dimension $d' < d$. Feature selection methods are usually designed to *improve* the classification by removing features that are too noisy or useless for the classification. We thus do not detail further these methods as they are not approximation algorithms *per se*. The interested reader will find some entries in the literature via references [25, 35, 60, 149]. Projection methods use random projections of the original points \mathbf{P} on spaces of dimension $d' \sim \log n$ in order to take advantage of the Johnson–Lindenstrauss lemma of norm conservation: the kernel computed from the projected features in dimension d' is thus an approximation of the true kernel with high probability. We refer to the works [64, 116] for more details.

5.4 Sampling Given the Similarity Graph

We now suppose that the similarity graph is either given (e.g., in cases where the original data *is* a graph) or has been well approximated (by approximate k -NN search, for instance) and concentrate on sampling-based methods that aim to reduce the cost of computing the first k eigenvectors of \mathbf{R} .

These methods predominantly aim to approximate \mathbf{R} by a smaller matrix $\tilde{\mathbf{R}}$ of size m . The eigen-decomposition is done in \mathbb{R}^m which can be significantly cheaper when $m \ll n$. In addition, each method comes with a fast way of lifting vectors from \mathbb{R}^m back to \mathbb{R}^n (this is usually a linear transformation). After lifting, the eigenvectors of $\tilde{\mathbf{R}}$ are used as a proxy for those of \mathbf{R} .

Unlike the previous section where a strong approximation guarantee of the exact embedding \mathbf{U}_k by an efficiently computed $\tilde{\mathbf{U}}_k$ was a distant and difficult goal to achieve in itself, we will see in this section that the knowledge of the similarity

¹³The recall rate for a node is the number of correctly identified k -NN divided by k . The recall rate for a k -NN graph is the average recall rate over all nodes.

graph not only enables to obtain such strong approximation guarantees, but also enables to control how the error on \mathbf{U}_k transfers as an error on the k -means cost.

To be more precise, recall (5.1) defining the k -means cost $f(\mathbf{C}; \mathbf{X})$ associated with the n points $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and a centroid set \mathbf{C} . Now, suppose that we have identified a set of n points $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1 | \dots | \tilde{\mathbf{x}}_n)$ that are meant to approximate the exact spectral embedding \mathbf{X} . Moreover, let \mathbf{C}^* (resp. $\tilde{\mathbf{C}}^*$) be the optimal set of k centroids minimizing the k -means cost on \mathbf{X} (resp. $\tilde{\mathbf{X}}$). We will see that several (not all) approximation methods of this section achieve an end-to-end approximation guarantee of the form

$$\left| f(\mathbf{C}^*; \mathbf{X})^{1/2} - f(\tilde{\mathbf{C}}^*; \mathbf{X})^{1/2} \right| \leq \epsilon,$$

for some small ϵ with—at least—constant probability. Such an end-to-end guarantee is indeed more desirable than a simple guarantee on the distance between \mathbf{U}_k and $\tilde{\mathbf{U}}_k$: it informs us on the approximation quality of the attained clustering.

5.4.1 Nyström-Based Methods

The Nyström-based methods discussed in Sect. 5.3.1 are also applicable here. Let us concentrate on the choice $\mathbf{R} = \mathbf{L}_n$ to illustrate the main ideas. As explained in Sect. 5.3.1, the tailing k eigenvectors \mathbf{U}_k of \mathbf{L}_n can be interpreted as the top k eigenvectors of the PSD matrix $\mathbf{A} = 2\mathbf{I} - \mathbf{L}_n$. As such, the span of the top- k eigenvectors of $\tilde{\mathbf{A}}_k$, $\text{span}(\tilde{\mathbf{U}}_k)$, obtained by running Algorithm 3 on \mathbf{A} should approximate the span of \mathbf{U}_k . Now, how does one go from Nyström theorems such as Theorem 5.2 to error bounds on the k -means cost function?

The first step towards an end-to-end guarantee relies on the following result:

Lemma 1 (See the Proof of Theorem 6 in [21]) *Denote by $\tilde{\mathbf{C}}^*$ the optimal centroid set obtained by solving k -means on the rows of $\tilde{\mathbf{U}}_k$. It holds that*

$$\left| f(\mathbf{C}^*; \mathbf{X})^{1/2} - f(\tilde{\mathbf{C}}^*; \mathbf{X})^{1/2} \right| \leq 2\|\mathbf{E}\|_F, \quad (5.16)$$

where $\mathbf{E} = \mathbf{U}_k \mathbf{U}_k^\top - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^\top$.

This means that the error made by considering the optimal k -means solution based on $\tilde{\mathbf{U}}_k$ (instead of \mathbf{U}_k) is controlled by the Frobenius norm of the projector difference $\mathbf{E} = \mathbf{U}_k \mathbf{U}_k^\top - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^\top$. Furthermore, since¹⁴ $\|\mathbf{E}\|_F \leq \sqrt{2k}\|\mathbf{E}\|_2$ and

¹⁴Based on three arguments: (i) for any two matrices \mathbf{M}_1 and \mathbf{M}_2 of rank r_1 and r_2 it holds that $\text{rank}(\mathbf{M}_1 + \mathbf{M}_2) \leq r_1 + r_2$, (ii) for any matrix \mathbf{M} of rank r , $\|\mathbf{M}\|_F \leq \sqrt{r}\|\mathbf{M}\|_2$, and (iii) both \mathbf{U}_k and $\tilde{\mathbf{U}}_k$ are of rank k .

$\|\mathbf{E}\|_2 = \|\sin(\Theta(\mathbf{U}_k, \tilde{\mathbf{U}}_k))\|_2$, we can apply the Davis–Kahan $\sin \Theta$ perturbation theorem (see, for instance, Section VII of [16]) and, provided that $\sigma_k - \tilde{\sigma}_{k+1} > 0$, obtain:

$$\|\mathbf{E}\|_F \leq \sqrt{2k}\|\mathbf{E}\|_2 \leq \sqrt{2k} \frac{\|\mathbf{A} - \tilde{\mathbf{A}}\|_2}{\sigma_k - \tilde{\sigma}_{k+1}},$$

where $\{\sigma_i\}$ (resp. $\{\tilde{\sigma}_i\}$) are the singular values of \mathbf{A} (resp. $\tilde{\mathbf{A}}$) ordered decreasingly.¹⁵ The final bound is obtained by combining the above with the leverage score sampling bound given by Theorem 5.2.

Theorem 5.4 *Let $\tilde{\mathbf{U}}_k$ be the eigenvectors obtained by running Algorithm 3 on $\mathbf{A} = 2\mathbf{I} - \mathbf{L}_n$ (with the leverage score sampling scheme for the m samples S of step 1). Denote by $\tilde{\mathbf{C}}^*$ the optimal centroid set obtained by solving k -means on the rows of $\tilde{\mathbf{U}}_k$. Then, for some constant $C > 1$, we have*

$$\left| f(\mathbf{C}^*; \mathbf{X})^{1/2} - f(\tilde{\mathbf{C}}^*; \mathbf{X})^{1/2} \right| \leq 2 \frac{\sqrt{2k}}{\sigma_k - \tilde{\sigma}_{k+1}} \left(\sigma_{k+1}(\mathbf{A}) + \frac{Ck \log(k/\delta)}{m} \sum_{j=k+1}^n \sigma_j \right)$$

with probability at least $0.8 - 2\delta$.

Examining the above bound one notices that $2\sqrt{2k} \frac{\sigma_{k+1}(\mathbf{A})}{\sigma_k - \tilde{\sigma}_{k+1}}$ is independent of the number of samples. The incompressibility of this error term emanates from \mathbf{A} being (in general) different from its best low-rank approximation. On the other hand, all remaining error terms can be made independent of k and n by setting

$$m = \mathcal{O} \left(k\sqrt{k} \log k \sum_{j=k+1}^n \frac{\sigma_j}{\sigma_k - \tilde{\sigma}_{k+1}} \right).$$

This end-to-end guarantee is not satisfactory for several reasons. First of all, it relies on the assumption $\sigma_k > \tilde{\sigma}_{k+1}$, which is not necessarily true. Moreover, the Davis–Kahan theorem could in theory guarantee $\|\mathbf{E}\|_2 \leq \|\mathbf{A}_k - \tilde{\mathbf{A}}_k\|_2/\sigma_k$ and $\|\mathbf{E}\|_2 \leq \|\mathbf{A} - \tilde{\mathbf{A}}_k\|_2/\sigma_k$, which are stronger than the bound depending on $\|\mathbf{A} - \tilde{\mathbf{A}}\|_2$ that we used. Unfortunately, Nyström approximation theorems do not give controls on $\|\mathbf{A}_k - \tilde{\mathbf{A}}_k\|_2$ nor on $\|\mathbf{A} - \tilde{\mathbf{A}}_k\|_2$, impeding tighter end-to-end bounds.

5.4.2 Graph Coarsening

Inspired by the algebraic multi-grid, researchers realized early on that a natural way to accelerate spectral clustering is by graph coarsening [38, 61, 69]. Here,

¹⁵Note that, in our setting, $\mathbf{A} = 2\mathbf{I} - \mathbf{L}_n$ and $\sigma_k = 2 - \lambda_k$.

instead of solving the clustering problem directly on G , one may first reduce it to a coarser graph G_c consisting of $m \ll n$ nodes using a multi-level graph coarsening procedure. The expensive eigen-decomposition computation is done at a lower cost on the representative matrix of the small graph and the final spectral embedding is obtained by inexpensively lifting and refining the result.

In the notation of [91], coarsening involves a sequence of $c + 1$ graphs

$$G = G_0 = (V_0, E_0, \mathbf{W}_0) \quad G_1 = (V_1, E_1, \mathbf{W}_1) \quad \cdots \quad G_c = (V_c, E_c, \mathbf{W}_c) \quad (5.17)$$

of decreasing size $n = n_0 > n_1 > \cdots > n_c = m$, where each vertex of G_ℓ represents one or more vertices of $G_{\ell-1}$. To express coarsening in algebraic form, we suppose that $\mathbf{L}(G_0) = \mathbf{L}$ is the combinatorial Laplacian associated with G . We then obtain $\mathbf{L}(G_c)$ by applying the following repeatedly:

$$\mathbf{L}(G_\ell) = \mathbf{P}_\ell^\mp \mathbf{L}(G_{\ell-1}) \mathbf{P}_\ell^+, \quad (5.18)$$

where $\mathbf{P}_\ell \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$ is a matrix with more columns than rows, $\ell = 1, 2, \dots, c$ is the level of the reduction, and symbol \mp denotes the transposed pseudo-inverse. An eigenvector $\tilde{\mathbf{u}} \in \mathbb{R}^m$ of $\mathbf{L}(G_c)$ is lifted back to \mathbb{R}^n by backwards recursion

$$\tilde{\mathbf{u}}_{\ell-1} = \mathbf{P}_\ell \tilde{\mathbf{u}}_\ell,$$

where $\tilde{\mathbf{u}}_c = \tilde{\mathbf{u}}$.

Matrices $\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_c$ are determined by the transformation performed at each level. Specifically, one should define for each level a surjective map $\varphi_\ell : V_{\ell-1} \rightarrow V_\ell$ between the original vertex set $V_{\ell-1}$ and the smaller vertex set V_ℓ . We refer to the set of vertices $V_{\ell-1}^{(r)} \subseteq V_{\ell-1}$ mapped onto the same vertex v'_r of V_ℓ as a *contraction set*:

$$V_{\ell-1}^{(r)} = \{v \in V_{\ell-1} : \varphi_\ell(v) = v'_r\}$$

It is easy to deduce from the above that contraction sets induce a partitioning of $V_{\ell-1}$ into n_ℓ subgraphs, each corresponding to a single vertex of V_ℓ .

Then, for any $v'_r \in V_\ell$ and $v_i \in V_{\ell-1}$, matrices $\mathbf{P}_\ell \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$ and $\mathbf{P}_\ell^+ \in \mathbb{R}^{n_{\ell-1} \times n_\ell}$ are given by:

$$\mathbf{P}_\ell(r, i) = \begin{cases} \frac{1}{|V_{\ell-1}^{(r)}|} & \text{if } v_i \in V_{\ell-1}^{(r)} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \mathbf{P}_\ell^+(i, r) = \begin{cases} 1 & \text{if } v_i \in V_{\ell-1}^{(r)} \\ 0 & \text{otherwise.} \end{cases}$$

The preceding construction is the only one that guarantees that every $\mathbf{L}(G_\ell)$ will be the combinatorial Laplacian associated with G_ℓ [90].

Note that from a computational perspective the reduction is very efficient and can be carried out in linear time: each coarsening level entails multiplication by a sparse

matrix, meaning that $\mathcal{O}(e)$ and $\mathcal{O}(n)$ operations suffice, respectively, to coarsen \mathbf{L} and lift any vector (such as the eigenvectors of $\mathbf{L}(G_c)$) from \mathbb{R}^m back to \mathbb{R}^n .

5.4.2.1 Coarsening for Spectral Clustering

Using coarsening effectively boils down to determining for each ℓ how to partition $G_{\ell-1}$ into n_ℓ contraction sets $V_\ell^{(1)}, \dots, V_\ell^{(n_\ell)}$, such that, after lifting, the first k eigenvectors $\tilde{\mathbf{U}}_k$ of $\mathbf{L}(G_c)$ approximate the spectral embedding \mathbf{U}_k derived from \mathbf{L} . Alternatively, one may also solve the k -means problem in the small dimension and only lift the resulting cluster assignments [38]. This scheme is computationally superior but we will not discuss it here as it does not come with any guarantees.

Perhaps the most *simple* (and common) method of forming contraction sets is by the *heavy-edge matching heuristic*—originally developed in the multi-grid literature and first considered for graph partitioning in [69]. This method is derived based on the intuition that the larger the weight of an edge, the less likely it will be that the vertices it connects will reside in different clusters. We should therefore aim to contract pairs of vertices connected by a heavy edge (i.e., of large weight) first. Let us consider this case further. By focusing on edges, we basically constrain ourselves by enforcing that every contraction set $V_\ell^{(r)}$ contains either two nodes connected by an edge or a single node, signifying that said node is chosen to remain as is in the coarser graph. As such, we can reformulate the problem of selecting contraction sets at each level as that of selecting the largest number of edges (to attain the largest reduction), while also striving to make the cumulative sum of selected edge weights as large as possible (giving preference to heavy edges). This is exactly the *maximum weight matching problem*, which can be approximated in linear time [44].

A plethora of numerical evidence motivates the use of matching-based coarsening methods, such as the heavy-edge heuristic, for accelerating spectral clustering [38, 69, 115]. From a theoretical perspective, the approximation quality of matching-based methods was characterized in [91]. Therein, the matching was constructed in the following randomized manner:

Algorithm 5. Randomized edge contraction (one level) [91]

Input. A graph $G = (V, E)$

1. Associate with each $e_{ij} \in E$ a probability $p_{ij} > 0$.
2. While $|E| > 0$:
 - (a) Draw a sample e_{ij} from E with probability $\propto p_{ij}$.
 - (b) Remove from E both e_{ij} as well as all edges sharing a common endpoint with it.
 - (c) Construct contraction set (v_i, v_j) .

Output: Contraction sets

The following approximation result is known:

Theorem 5.5 (Corollary 5.1 in [91]) *Consider a graph with bounded degrees $d_i \ll n$ and $\lambda_k \leq \min_{e_{ij} \in E} \left\{ \frac{d_i + d_j}{2} \right\}$. Suppose that the graph is coarsened by Algorithm 5, using a heavy-edge potential such that $p_{ij} \propto w_{ij}$. For sufficiently large n , a single level, and $\delta > 0$,*

$$\left| f(\mathbf{C}^*; \mathbf{X})^{1/2} - f(\tilde{\mathbf{C}}^*; \mathbf{X})^{1/2} \right| = \mathcal{O} \left(\sqrt{\frac{1 - \frac{m}{n} \sum_{\ell=2}^k \lambda_\ell}{\delta \lambda_{k+1} - \lambda_k}} \right)$$

with probability at least $1 - \delta$. Above, $\tilde{\mathbf{C}}^$ is the optimal k -means solution when using the lifted eigenvectors of \mathbf{L}_c as a spectral embedding.*

We deduce that coarsening works better when the spectral clustering problem is easy (as quantified by the weighted gap $\sum_{\ell=2}^k \lambda_\ell / (\lambda_{k+1} - \lambda_k)$) and the achieved error is linear on the reduction ratio $1 - m/n$.

There also exist more advanced techniques for selecting contraction sets that come with stronger guarantees w.r.t. the attained reduction and quality of approximation, but feature running time that is not smaller than that of spectral clustering [90]. In particular, these work also with the normalized Laplacian and can be used to achieve multi-level reduction. Roughly, their strategy is to identify and contract sets $\mathbf{S} \subset V$ for which $\mathbf{x}(i) \approx \mathbf{x}(j)$ for all vectors $\mathbf{x} \in \mathbf{U}_k$ and $v_i, v_j \in \mathbf{S}$. This strategy ensures that the best partitionings of G are preserved by coarsening. We will not expand on these methods here as they do not aim to improve the running time of spectral clustering.

5.4.3 Other Approaches

In the following, we present two additional approaches for approximately computing spectral embeddings. The former can be interpreted as a sampling-based method (but in a different manner than the techniques discussed so far), whereas the latter is only vaguely linked to sampling. Nevertheless, we find that both techniques are very interesting and merit a brief discussion.

5.4.3.1 Spectral Sparsification

This approach is best suited for cases when the input of spectral clustering is directly a graph.¹⁶ Different from the methods discussed earlier, here the aim is to identify

¹⁶When one starts from a set of points, it is preferable to sparsify the graph by retaining a constant number of nearest neighbors for each point. The resulting nearest neighbor graph has already $O(n)$ edges, which is the smallest possible.

a Laplacian matrix $\tilde{\mathbf{L}}$ of the same size as \mathbf{L} but with fewer entries. Additionally, it should be ensured that

$$\frac{1}{1+\epsilon} \mathbf{x}^\top \mathbf{L} \mathbf{x} \leq \mathbf{x}^\top \tilde{\mathbf{L}} \mathbf{x} \leq (1+\epsilon) \mathbf{x}^\top \mathbf{L} \mathbf{x} \quad \text{for all } \mathbf{x} \in \mathbb{R}^n \quad (5.19)$$

for some small constant $\epsilon > 0$ [125]. Most fast algorithms for spectral sparsification entail sampling $\mathcal{O}(n \log n)$ edges from the total edges present in the graph. Different sampling schemes are possible [72, 124], but the most popular ones entail sampling edges with replacement based on their effective resistance. It should be noted that though computing all effective resistances exactly can be computationally prohibitive, the effective resistance of edges can be approximated in nearly linear time on the number of edges based on a Johnson–Lindenstrauss argument [124].

There are different ways to use sparsification in order to accelerate spectral clustering. The most direct one is to exploit the fact that the eigenvalues $\tilde{\lambda}_k$ and eigenvectors $\tilde{\mathbf{U}}_k$ of $\tilde{\mathbf{L}}$ approximate, respectively, the eigenvalues and eigenvectors of \mathbf{L} up to multiplicative error. This yields the same flavor of guarantees as in graph coarsening and ensures that the computational complexity of the partial eigen-decomposition will decrease when $e = \omega(n \log n)$. A variation of this idea was considered in [140], though the latter did not provide a complete error and complexity analysis. Alternative approaches are also possible. We refer the interested reader to [136] for a rigorous argument that invokes a Laplacian solver.

Despite these exciting developments, we should mention that the overwhelming majority of graph sparsification algorithms remain in the realm of theory. That is, we are currently not aware of any practical and competitive implementation and thus retain a measure of skepticism with regard to their utility in the setting of spectral clustering.

5.4.3.2 Random Eigenspace Projection

There also exist approaches that do not explicitly rely on sampling. The key starting point here is that, with regard to spectral clustering, one does not need the eigenvectors exactly—any rotation of \mathbf{U}_k suffices (indeed, k -means is an algorithm based on distances and rotations conserve distances). Even more generally, consider $\tilde{\mathbf{U}}_k \in \mathbb{R}^{n \times m}$ with $m \geq k$ and denote:

$$\epsilon = \min_{\mathbf{Q} \in \mathcal{Q}} \|\mathbf{U}_k \mathbf{I}_{k \times m} \mathbf{Q} - \tilde{\mathbf{U}}_k\|_F,$$

where \mathcal{Q} is the space of $m \times m$ unitary matrices and $\mathbf{I}_{k \times m}$ consists of the first k rows of an $m \times m$ identity matrix.

The following lemma (which is a generalization of Lemma 1) shows how ϵ can be used to provide control on the k -means error:

Lemma 2 (Lemma 3.1 in [95])

Let $\tilde{\mathbf{C}}^*$ be the optimal solution of the k -means problem on $\tilde{\mathbf{U}}_k$. It holds that¹⁷

$$\left| f(\mathbf{C}^*; \mathbf{X})^{1/2} - f(\tilde{\mathbf{C}}^*; \mathbf{X})^{1/2} \right| \leq 2\epsilon. \quad (5.20)$$

There exists (at least) two approaches to efficiently compute $\tilde{\mathbf{U}}_k$ while controlling ϵ [21, 130] (see also related work in [58]). We will consider here a simple variant of the one proposed in [130] and further analyzed in [95]. Let $\mathbf{G} \in \mathbb{R}^{n \times m}$ be a random Gaussian matrix with centered i.i.d. entries, each having variance $\frac{1}{m}$. Furthermore, suppose that we project \mathbf{G} onto $\text{span}(\mathbf{U}_k)$ by multiplying each one of its columns by an ideal projector \mathbf{P}_k defined as

$$\mathbf{P}_k = \mathbf{U} \begin{pmatrix} \mathbf{I}_k & 0 \\ 0 & 0 \end{pmatrix} \mathbf{U}^\top. \quad (5.21)$$

Theorem 5.6 ([95, 130]) Let $\tilde{\mathbf{C}}^*$ be the optimal solution of the k -means problem on the rows of $\tilde{\mathbf{U}}_k = \mathbf{P}_k \mathbf{G}$. For every $\delta \geq 0$, one has

$$\left| f(\mathbf{C}^*; \mathbf{X})^{1/2} - f(\tilde{\mathbf{C}}^*; \mathbf{X})^{1/2} \right| \leq 2\sqrt{\frac{k}{m}} (\sqrt{k} + \delta), \quad (5.22)$$

with probability at least $1 - \exp(-\delta^2/2)$.

This result means that for an ideal projector \mathbf{P}_k , dimension $m = \mathcal{O}(k^2)$ suffices to guarantee good approximation (since the error becomes independent of k and n)! A similar argument also holds when the entries of \mathbf{G} , instead of being Gaussian, are selected i.i.d. from $\{-\sqrt{3}, 0, +\sqrt{3}\}$ with probabilities $\{1/6, 2/3, 1/6\}$, respectively [1]. This construction has the benefit of being sparser and, moreover, is reminiscent of sampling. It should be noted that in [130], $m = \mathcal{O}(\log n)$ was deemed enough because one only wanted that the distance between two rows of \mathbf{U}_k

¹⁷**A remark on the definition of the k -means cost.** Note that, here, the lines $\tilde{\mathbf{X}}$ of $\tilde{\mathbf{U}}_k$ are points in dimension $m \geq k$, such that the optimal centroid set $\tilde{\mathbf{C}}^*$ minimizing the k -means cost on $\tilde{\mathbf{X}}$ is a set of k points in dimension $m \geq k$. In this context, the notation $f(\tilde{\mathbf{C}}^*; \mathbf{X})$ is ill-defined: it is a sum of distances between points that do not necessarily have the same dimension. We abuse notations and give the following meaning to $f(\tilde{\mathbf{C}}; \mathbf{X})$. First, consider the matrix form of the k -means cost, as used in the proofs of Lemmas 1 and 2: $f(\tilde{\mathbf{C}}; \mathbf{X}) = \|\mathbf{X} - \mathbf{C}\tilde{\mathbf{C}}^\top \mathbf{X}\|_F^2$, where $\mathbf{X} = (\mathbf{x}_1 | \dots | \mathbf{x}_n)^\top \in \mathbb{R}^{n \times k}$ and $\mathbf{C} \in \mathbb{R}^{n \times k}$ is the (weighted) cluster indicator matrix associated to the Voronoi tessellation of \mathbf{X} given $\tilde{\mathbf{C}}$: $\mathbf{C}_{i\ell} = 1/\sqrt{s_\ell}$ if data point i belongs to cluster ℓ , and 0 otherwise, where s_ℓ is the size of cluster ℓ . Now, let $\tilde{\mathbf{C}} \in \mathbb{R}^{n \times k}$ be the cluster indicator matrix associated with the Voronoi tessellation of $\tilde{\mathbf{X}}$ given $\tilde{\mathbf{C}}$. One writes: $f(\tilde{\mathbf{C}}; \mathbf{X}) = \|\mathbf{X} - \tilde{\mathbf{C}}\tilde{\mathbf{C}}^\top \mathbf{X}\|_F^2$.

was approximated by the distance between the same two rows of $\tilde{\mathbf{U}}_k$. There was in fact no end-to-end control on the k -means error.

The discussion so far assumed that \mathbf{P}_k is an ideal projector onto $\text{span}(\mathbf{U}_k)$. However, in practice one does not have access to this projector as we are in fact in the process of *computing* \mathbf{U}_k . One may choose to approximate the action of \mathbf{P}_k by an application of a matrix function h on the representative matrix \mathbf{R} [111, 129]. Assuming a point λ_* in the interval $[\lambda_k, \lambda_{k+1})$ is known, one may select a polynomial [122] or rational function [65, 92] that approximates the ideal low-pass response, i.e., $h(\lambda) = 1$ if $\lambda \leq \lambda_*$ and $h(\lambda) = 0$, otherwise. The approximated projector $\tilde{\mathbf{P}}_k = h(\mathbf{R})$ can be designed to be very close to \mathbf{P}_k . For instance, in the case of Chebyshev polynomials of order c using the arguments of [80, Lemma 1] it is easy to prove that w.h.p. using $h(\mathbf{R})$ instead of \mathbf{P}_k does not add more than $\mathcal{O}(c^{-c}\sqrt{n})$ error in (5.20). Furthermore, the operation $\tilde{\mathbf{P}}_k\mathbf{G}$ can conveniently be computed in $\mathcal{O}(mce)$ time via this polynomial approximation.

The last ingredient needed for this approximation is λ_* , i.e., a point in the interval $[\lambda_k, \lambda_{k+1})$. Finding efficiently a valid λ_* is difficult. An option is to rely on eigencount techniques [39, 105, 109] to find one in¹⁸ $\mathcal{O}(ck^2(\log n)(e + n \log(\lambda_n/(\lambda_{k+1} - \lambda_k))))$ time, which features similar complexity as the Lanczos method (see discussion in Sect. 5.2.3). Another option is to content oneself with values of λ_* known only to be close to the interval $[\lambda_k, \lambda_{k+1})$, but thereby loosing the end-to-end guarantee [130].

5.5 Sampling in the Spectral Feature Space

Having computed (or approximated) the spectral embedding $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, what remains is to solve the k -means problem on \mathbf{X} , in order to obtain k centroids together with the associated k classes obtained after Voronoi tessellation.

The usual heuristic used to solve the k -means problem, namely the Lloyd-Max algorithm, is already very efficient as it runs in $\mathcal{O}(nk^2t)$ time as seen in Sect. 5.2.3. Nonetheless, this section considers ways to accelerate k -means even further. In the following, we classify the relevant literature in five categories and point towards representative references for each case. In our effort to provide depth (as well as

¹⁸*Proof sketch:* Given $\lambda \in (0, \lambda_n)$, denote by j the largest integer such that $\lambda_j \leq \lambda$ and by \mathbf{P}_j the orthogonal projector on \mathbf{U}_j . Let $\mathbf{G} \in \mathbb{R}^{n \times m'}$ be a random Gaussian matrix with centered i.i.d. entries, each having variance $\frac{1}{m'}$ and denote by $\hat{j} = \|\mathbf{P}_j\mathbf{G}\|_F$. Relying on Theorem 4.1 (and the following discussion in Section 4.2) of [109] with $\mathbf{E}_\lambda = \mathbf{0}$, one has with prob. at least $1 - \epsilon$ that $(1 - \delta)j \leq \hat{j} \leq (1 + \delta)j$ for all $j = 1, \dots, n$ provided $m' \geq \frac{1}{\delta^2} \log \frac{n}{\epsilon}$. Setting $\delta = 1/(2k + 3)$ gives w.h.p. that $\frac{2k+2}{2k+3}j \leq \hat{j} \leq \frac{2k+4}{2k+3}j$ for all $j = 1, \dots, n$ provided $m' \geq \mathcal{O}(k^2 \log n)$. This implies that w.h.p. for every $j \leq k + 1$ it must be that $\text{round}(\hat{j}) = j$, whereas when $j > k + 1$ we have $\text{round}(\hat{j}) > k + 1$. Note that $\text{round}(\hat{j})$ is the closest integer to \hat{j} . By dichotomy on $\lambda \in (0, \lambda_n)$, one thus finds a λ_* in time $\mathcal{O}(ck^2(\log n)(e + n \log(\lambda_n/(\lambda_{k+1} - \lambda_k))))$.

breadth) of presentation, the rest of the section details only methods that belong to the first and last categories.

- **Exact acceleration of Lloyd-Max.** There exists exact accelerated Lloyd-Max algorithms, some of them based on avoiding unnecessary distance calculations using the triangular inequality [59, 101], or on optimized data organization [67], and others concentrating on clever initializations [6, 104]. Unlike the former methods, the latter methods involve sampling and are discussed in Sect. 5.5.1.
- **Approximate acceleration of Lloyd-Max.** Approximately accelerating the Lloyd-Max algorithm has also received attention, for instance, via approximate nearest neighbor methods [108], via cluster closure [142], or via applying Lloyd-Max hierarchically (in the large k context) [103]. An approach involving sampling is introduced in [119]: it is based on mini-batches sampled uniformly at random from \mathbf{X} . We will not discuss further this method as it does not come with guarantees on the cost of the obtained solution.
- **Methods involving sampling in the Fourier domain.** There are a few sampling-based heuristics to solve the k -means problem that are different from the Lloyd-Max algorithm. For instance, the work in [70] proposes to sample in the frequency domain to obtain a sketch from which one may recover the centroids with an orthogonal matching pursuit algorithm specifically tailored to this kind of compressive learning task [55]. These methods are reminiscent of the random Fourier features sketching approach introduced in Sect. 5.3.2. We will not discuss them further.
- **Methods involving sampling features.** Similarly to ideas presented in Sect. 5.3.4.2 but here specific to the k -means setting, some works reduce the ambient dimension of the vectors, either by selecting a limited number of features [3, 20] or by embedding all points in a lower dimension using random projections [22, 33, 93]. The tightest results to day are a $(1 + \epsilon)$ multiplicative error on the k -means cost f either by randomly selecting $\mathcal{O}(\epsilon^{-2}k \log k)$ features or by projecting them on a random space of dimension $\mathcal{O}(\epsilon^{-2} \log(k/\epsilon))$ (sublinear in $k!$). The sampling result is useless in the spectral clustering setting as the ambient dimension of the spectral features is already k . The projection result could in principle be applied in our setting, to reduce the cost of the k -means step to $\mathcal{O}(tnk \log k)$. We will nevertheless not discuss it further in this chapter.
- **Methods involving sampling points.** Finally, the last group of existing methods are the ones that solve k -means on a subset \mathbf{S} of \mathbf{X} , before lifting back the result on the whole dataset. We classify such methods in two categories. In Sect. 5.5.2, we detail methods that are graph-agnostic, meaning that they apply to *any* k -means problem; and in Sect. 5.5.3 we discuss methods that explicitly rely on the fact that the features \mathbf{x} were in fact obtained from a known graph. We argue that the latter are better suited to the spectral clustering problem.

5.5.1 Clever Initialization of the Lloyd-Max Algorithm

Recall that the k -means objective on \mathbf{X} is to find the k centroids $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_k)$ that minimize the following cost function:

$$f(\mathbf{C}; \mathbf{X}) = \sum_{\mathbf{x} \in \mathbf{X}} \min_{\mathbf{c} \in \mathbf{C}} \|\mathbf{x} - \mathbf{c}\|_2^2. \quad (5.23)$$

and that $\mathbf{C}^* = \arg \min_{\mathbf{C}} f(\mathbf{C}; \mathbf{X})$ is the optimal solution attaining cost $f^* = f(\mathbf{C}^*; \mathbf{X})$. Recall also that the Lloyd-Max algorithm (see Algorithm 2) converges to a local minimum of f that we will denote by \mathbf{C}_{lm} , for which the cost function equals $f_{\text{lm}} = f(\mathbf{C}_{\text{lm}}; \mathbf{X})$. It is crucial to note that the initialization of centroids \mathbf{C}_{ini} in the first step of the Lloyd-Max algorithm, which usually is done by randomly selecting k points in \mathbf{X} , is what determines the distance $|f^* - f_{\text{lm}}|$ to the optimal. As such, significant efforts have been devoted to smartly selecting \mathbf{C}_{ini} by various sampling schemes.

As usual, we also face here the usual trade-off between sampling *effectively* and *efficiently*. The fastest sampling method is of course uniformly at random, but it does not come with any guarantee on the quality of the local minimum \mathbf{C}_{lm} it leads to. An alternative sampling scheme, called k -means++ initialization, is based on the following more general \mathbf{D}^2 -sampling algorithm.

Algorithm 6: \mathbf{D}^2 -sampling

Input: \mathbf{X} , m the number of required samples

1. Initialize \mathbf{B} with any \mathbf{x} chosen uniformly at random from \mathbf{X} .
2. Iterate the following steps until \mathbf{B} contains m elements:
 - (a) Compute $d_i = \min_{b \in \mathbf{B}} \|\mathbf{x}_i - b\|_2^2$.
 - (b) Define the probability of sampling \mathbf{x}_i as $d_i / \sum_i d_i$.
 - (c) Sample \mathbf{x}_{new} from this probability distribution and add it to \mathbf{B} .

Output: \mathbf{B} a sample set of size m .

k -means++ initialization boils down to running Algorithm 6 with $m = k$ to obtain a set of k initial centroids. Importantly, when the Lloyd-Max heuristic is run with this initialization, the following guarantee holds:

Theorem 5.7 ([6]) *For any set of data points, the cost f_{lm} obtained after Lloyd-Max initialized with k -means++ is controlled in expectation: $\mathbb{E}(f_{\text{lm}}) \leq 8(\log k + 2)f^*$.*

In terms of computation cost, \mathbf{D}^2 -sampling with $m = k$ runs in $\mathcal{O}(nkd)$, that is, $\mathcal{O}(nk^2)$ in our setting of a spectral embedding \mathbf{X} in dimension k . This work inspired

other initialization techniques that come with similar guarantees and are in some cases faster [10, 12]. The interested reader is referred to the review [27] for further analyses on the initialization of k -means.

5.5.2 Graph Agnostic Sampling Methods: Coresets

The rest of Sect. 5.5 considers sampling methods that fall in the following framework: (i) sample a subset \mathbf{S} of \mathbf{X} , (ii) solve k -means on \mathbf{S} , (iii) lift the result back on the whole dataset \mathbf{X} . Section 5.5.2 focuses on coresets: general sampling methods designed for any arbitrary k -means problem, whereas in Sect. 5.5.3, we will take into account the specific nature of the spectral features encountered in spectral clustering algorithms.

5.5.2.1 Definition

Let $\mathbf{S} \subset \mathbf{X}$ be a subset of \mathbf{X} of size m . To each element $s \in \mathbf{S}$ associate a weight $\omega(s) \in \mathbb{R}^+$. Define the estimated k -means cost associated with the weighted set \mathbf{S} as:

$$\tilde{f}(\mathbf{C}; \mathbf{S}) = \sum_{s \in \mathbf{S}} \omega(s) \min_{\mathbf{c} \in \mathbf{C}} \|s - \mathbf{c}\|_2^2. \quad (5.24)$$

Definition 5.2 (Coreset) Let $\epsilon \in (0, \frac{1}{2})$. The weighted subset \mathbf{S} is a ϵ -coreset for f on \mathbf{X} if, for every set \mathbf{C} , the estimated cost is equal to the exact cost up to a relative error:

$$\forall \mathbf{C} \quad \left| \frac{\tilde{f}(\mathbf{C}; \mathbf{S})}{f(\mathbf{C}; \mathbf{X})} - 1 \right| \leq \epsilon. \quad (5.25)$$

This is the so-called strong coreset definition,¹⁹ as the ϵ -approximation is required for all \mathbf{C} . The great interest of finding a coreset \mathbf{S} comes from the following fact. Writing $\tilde{\mathbf{C}}^*$ the set minimizing \tilde{f} , the following inequalities hold:

$$(1 - \epsilon)f(\mathbf{C}^*; \mathbf{X}) \leq (1 - \epsilon)f(\tilde{\mathbf{C}}^*; \mathbf{X}) \leq \tilde{f}(\tilde{\mathbf{C}}^*; \mathbf{S}) \leq \tilde{f}(\mathbf{C}^*; \mathbf{S}) \leq (1 + \epsilon)f(\mathbf{C}^*; \mathbf{X}).$$

¹⁹A weaker version of this definition exists in the literature where the ϵ -approximation is only required for \mathbf{C}^* .

The first inequality comes from the fact that \mathbf{C}^* is optimal for f , the second and last inequality are justified by the coreset property of \mathbf{S} , and the third inequality comes from the optimality of $\tilde{\mathbf{C}}^*$ for \tilde{f} . This has two consequences:

1. First of all, since $\epsilon < \frac{1}{2}$:

$$f(\mathbf{C}^*; \mathbf{X}) \leq f(\tilde{\mathbf{C}}^*; \mathbf{X}) \leq (1 + 4\epsilon)f(\mathbf{C}^*; \mathbf{X}),$$

meaning that $\tilde{\mathbf{C}}^*$ is a well-controlled approximation of \mathbf{C}^* with a multiplicative error on the cost.

2. Estimating $\tilde{\mathbf{C}}^*$ can be done using the Lloyd-Max algorithm on the weighted subset²⁰ \mathbf{S} , thus reducing the computation time from $\mathcal{O}(nk^2)$ to $\mathcal{O}(mk^2)$.

Coreset methods for k -means thus follow the general procedure:

Algorithm 7. Coresets to avoid k -means on X

Input. X , sampling set size m , and number of clusters $k \leq m$.

1. Compute a weighted coreset \mathbf{S} of size m using a coreset-sampling algorithm.
2. Run the Lloyd-Max algorithm on the weighted set \mathbf{S} to obtain the set of k centroids $\tilde{\mathbf{C}}$.
3. “Closest-centroid lifting”: classify the whole dataset X based on the Voronoi cells of $\tilde{\mathbf{C}}$.

Output: A set of k centroids $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_k)$.

Coreset methods compete with one another on essentially two levels: the coreset size m should be as small as possible in order to decrease the time of Lloyd-Max on \mathbf{S} , and the coreset itself should be sampled efficiently (at least faster than running k -means on the whole dataset!), which turns out in fact to be a strong requirement. The reader interested in an overview of coreset construction techniques is referred to the recent review [99], as well as Chap. 2 of this book.

5.5.2.2 An Instance of Coreset-Sampling Algorithm

We focus on a particular coreset algorithm proposed in [11] that builds upon results developed in [45, 79]: it is not state of the art in terms of coreset size, but has the advantage of being easy to implement and fast enough to compute. It reads:

²⁰Generalizing Algorithm 2 to a weighted set is straightforward: in step 2b, instead of computing the center of each cluster, compute the weighted barycenter.

Algorithm 8: a coreset sampling algorithm [11]**Input.** \mathbf{X} , m the number of required samples, t an iteration number

1. Repeat t times: draw a set of size k using \mathbf{D}^2 -sampling. Out of the t sets obtained, keep the set \mathbf{B} that minimizes $f(\mathbf{B}; \mathbf{X})$.
2. $\alpha \leftarrow 16(\log k + 2)$
3. For each $b_\ell \in \mathbf{B}$, define \mathbf{B}_ℓ the set of points in \mathbf{X} in the Voronoi cell of b_ℓ
4. Set $\phi = \frac{1}{n} f(\mathbf{B}; \mathbf{X})$.
5. For each $b_\ell \in \mathbf{B}$ and each $\mathbf{x} \in \mathbf{B}_\ell$, define

$$s(\mathbf{x}) = \frac{\alpha}{\phi} \|\mathbf{x} - b_\ell\|_2^2 + \frac{2\alpha}{\phi |\mathbf{B}_\ell|} \sum_{\mathbf{x}' \in \mathbf{B}_\ell} \|\mathbf{x}' - b_\ell\|_2^2 + \frac{4n}{|\mathbf{B}_\ell|}$$

6. Define the probability of sampling \mathbf{x}_i as $\mathbf{p}_i = s(\mathbf{x}_i) / \sum_{\mathbf{x}} s(\mathbf{x})$
7. $\mathbf{S} \leftarrow$ sample m nodes i.i.d. with replacement from \mathbf{p} and associate to each sample s the weight $\omega_s = \frac{1}{m p_s}$.

Output: A weighted set \mathbf{S} of size m .

Theorem 2.5 of [11] states:

Theorem 5.8 *Let $\epsilon \in (0, 1/4)$ and $\delta \in (0, 1)$. Let \mathbf{S} be the output of Algorithm 8 with $t = \mathcal{O}(\log 1/\delta)$. Then, with probability at least $1 - \delta$, \mathbf{S} is a ϵ -coreset provided that:*

$$m = \Omega \left(\frac{k^4 \log k + k^2 \log 1/\delta}{\epsilon^2} \right). \quad (5.26)$$

The computation cost of running this coreset-sampling algorithm, running Lloyd-Max on the weighted coreset, and lifting the result back to \mathbf{X} is dominated, when²¹ $n \gg k$, by step 1 of Algorithm 6 and thus sums up to $\mathcal{O}(nk^2 \log 1/\delta)$.

Remark 1 The coreset-sampling strategy underlying this algorithm relies on the concept of sensitivity [79]. Many other constructions of coresets for k -means are possible [99] with better theoretical bounds than (5.26). Nevertheless, as the coreset line of research has been essentially theoretical, practical implementations of coreset-sampling algorithms are scarce. A notable exception is, for instance, the work in [49] that proposes a scalable hybrid coreset-inspired algorithm for k -means. Other exceptions are the sampling algorithms based on the farthest-first procedure, a variant of \mathbf{D}^2 -sampling that chooses each new sample to be $\arg \max_i d_i$ instead of drawing it according to a probability proportional to d_i . Once \mathbf{S} of size m is drawn,

²¹To be precise, the statement holds if $n \geq \mathcal{O} \left(\frac{k^4 \log k}{\epsilon^2 \log 1/\delta} \right)$.

then $\forall s \in \mathbf{S}$, each weight ω_s is set to be the cardinal of the Voronoi cell associated with s . Authors in [113] show that such weighted sets computed by different variants of the farthest-first algorithm are ϵ -coresets, but for values of ϵ that can be very large. For a fixed ϵ , the number of samples necessary to have a ϵ -coreset with this type of algorithm is unknown (see also Chap. 3 of this book).

5.5.3 Graph-Based Sampling Methods

The methods discussed so far in this section are graph agnostic both for the sampling procedure *and* the lifting: they do not take into account that, in spectral clustering, \mathbf{X} are in fact spectral features of a known graph.

A recent line of work [52, 53, 95, 130] based on graph signal processing (GSP) [118, 123] leverages this additional knowledge for accelerating both the sampling and the lifting steps. For the purpose of the following discussion, denote by $\mathbf{z}_\ell \in \mathbb{R}^n$ the ground truth indicator vector of cluster ℓ , i.e., $\mathbf{z}_\ell(i) = 1$ if node v_i is in cluster ℓ , and 0 otherwise. The goal of spectral clustering is, of course, to recover $\{\mathbf{z}_\ell\}_{\ell=1,\dots,k}$.

Broadly, GSP-based methods can be summarized in the following general methodology [130]:

Algorithm 9. Graph-based sampling strategies to avoid k -means on \mathbf{X}

Input. \mathbf{X} , m the number of required samples, k the number of desired clusters

1. Choose the random sampling strategy. Either:
 - (a) **uniform (i.i.d.)** Draw m i.i.d. samples uniformly.
 - (b) **leverage score (i.i.d.)** Compute $\forall \mathbf{x}_i, p_i^* = \|\mathbf{U}_k^\top \delta_i\|_2^2/k$. Draw m i.i.d. samples from \mathbf{p}^* . (optional:) set the weight of each sample s to $1/p_s^*$.
 - (c) **DPP** Sample a few times independently from a DPP with kernel $\mathbf{K}_k = \mathbf{U}_k \mathbf{U}_k^\top$. (optional:) set the weight of each sample s to $1/\pi_s$.
2. Run the Lloyd-Max algorithm on the (possibly weighted) set \mathbf{S} to obtain the k reduced cluster indicator vectors $\mathbf{z}'_\ell \in \mathbb{R}^m$.
3. Lift each reduced indicator vector $\{\mathbf{z}'_\ell\}_{\ell=1,\dots,k}$ to the full graph either with
 - (a) **Least-square** Solve (5.33) with $\mathbf{y} \leftarrow \mathbf{z}'_\ell$.
 - (b) **Tikhonov** Solve (5.34) with $\mathbf{y} \leftarrow \mathbf{z}'_\ell$.

(continued)

In both cases, \mathbf{P}_S should be set to $\frac{1}{N}\mathbf{I}_m$ if uniform sampling was chosen, to $\text{diag}(p_{s_1}^*, \dots, p_{s_m}^*)$ if leverage score sampling was chosen, and to $\text{diag}(\pi_{s_1}, \dots, \pi_{s_m})$ if DPP sampling was chosen.

4. Assign each node j to the cluster ℓ for which $\hat{\mathbf{z}}_\ell(j)/\|\hat{\mathbf{z}}_\ell\|_2$ is maximal.

Output: A partition of \mathbf{X} in k clusters

To aid understanding, let us start by a high-level description of Algorithm 9. The indicator vectors \mathbf{z}_ℓ are interpreted as graph signals that are (approximately) bandlimited on the similarity graph G (see Sect. 5.5.3.1 for a precise definition). As such, there is no need to measure these indicator vectors everywhere: one can take advantage of generalized Shannon-type sampling theorems to select the set \mathbf{S} of m nodes to measure (step 1). Then k -means is performed on \mathbf{S} to obtain the indicator vectors $\mathbf{z}_\ell^r \in \mathbb{R}^m$ on the sample set \mathbf{S} (step 2). These reduced indicator vectors are interpreted as noisy measurements of the global cluster indicator vectors \mathbf{z}_ℓ on \mathbf{S} . The solutions \mathbf{z}_ℓ^r are lifted back to \mathbf{X} as $\hat{\mathbf{z}}_\ell$ via solving an inverse problem taking into account the bandlimitedness assumption or via label-propagation on the graph structure reminiscent of semi-supervised learning techniques (step 3). As the lifted solutions $\hat{\mathbf{z}}_\ell$ do not have a binary structure as true indicator vectors should have, an additional assignment step is necessary: assign each node j to the class ℓ for which $\frac{\hat{\mathbf{z}}_\ell(j)}{\|\hat{\mathbf{z}}_\ell\|_2}$ is maximal (step 4).

The rest of this section is devoted to the discussion of the three sampling schemes as well as the two lifting procedures considered in this framework. To this end, we will first introduce a few graph signal processing (GSP) concepts in Sec. 5.5.3.1 before discussing in Sec. 5.5.3.2 several examples of graph sampling theorems appropriate to the spectral clustering context.

5.5.3.1 A Brief Introduction to Graph Signal Processing (GSP)

Denote by $\mathbf{U} = (\mathbf{u}_1 | \dots | \mathbf{u}_n) \in \mathbb{R}^{n \times n}$ the matrix of orthonormal eigenvectors of the Laplacian matrix \mathbf{L} , with the columns ordered according to their associated sorted eigenvalues: $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. In the GSP literature [118, 123], these eigenvectors are interpreted as graph Fourier modes for two main reasons:

- By analogy to the ring graph, whose Laplacian matrix is exactly the (symmetric) double derivative discrete operator, and is thus diagonal in the basis formed by the classical 1D discrete Fourier modes.
- A variational argument stemming from the Dirichlet form can be exploited to express eigenvectors \mathbf{u}_i of \mathbf{L} as the basis of minimal variation $\mathbf{x}^\top \mathbf{L} \mathbf{x} = \frac{1}{2} \sum_{ij} \mathbf{W}_{ij} [\mathbf{x}(i) - \mathbf{x}(j)]^2$ on G and eigenvalues λ_i as a sum of local variations of \mathbf{u}_i , i.e., a generalized graph frequency.

A *graph signal* $\mathbf{z} \in \mathbb{R}^n$ is a signal that is defined on the nodes of a graph: its i -th element is associated with node v_i . Given the previous discussion, the graph Fourier transform of \mathbf{z} , denoted by $\tilde{\mathbf{z}}$, is its projection on the graph Fourier modes: $\tilde{\mathbf{z}} = \mathbf{U}^\top \mathbf{z} \in \mathbb{R}^n$. The notion of graph filtering naturally follows as a multiplication in the Fourier domain. More precisely, define a real-valued filter function $h(\lambda)$ defined on $[0, \lambda_n]$. The signal \mathbf{x} filtered by h reads $\mathbf{U}h(\Lambda)\mathbf{U}^\top \mathbf{x}$, where we use the convention $h(\Lambda) = \text{diag}(h(\lambda_1), h(\lambda_2), \dots, h(\lambda_n))$. In the following, we will use the following notation for graph filter operators:

$$h(\mathbf{L}) = \mathbf{U}h(\Lambda)\mathbf{U}^\top. \quad (5.27)$$

For more details on the graph Fourier transform and filtering, their various definitions and interpretations, we refer the reader to [133].

Of interest for the discussion in this chapter, one may define bandlimited graph signals as linear combinations of the first few low-frequency Fourier modes. Writing $\mathbf{U}_k = (\mathbf{u}_1 | \dots | \mathbf{u}_k) \in \mathbb{R}^{n \times k}$, we have the formal definition:

Definition 5.3 (k -Bandlimited Graph Signal) A graph signal $\mathbf{z} \in \mathbb{R}^n$ is k -bandlimited if $\mathbf{z} \in \text{span}(\mathbf{U}_k)$, i.e., $\exists \boldsymbol{\alpha} \in \mathbb{R}^k$ such that $\mathbf{z} = \mathbf{U}_k \boldsymbol{\alpha}$.

To grasp why the notion of k -bandlimitedness lends itself naturally to the approximation of spectral clustering, consider momentarily a graph with k disconnected components and $\mathbf{z}_\ell \in \mathbb{R}^n$ the indicator vector of cluster ℓ . It is a well-known property of the (combinatorial) Laplacian that $\{\mathbf{z}_\ell\}_{\ell=1, \dots, k}$ form a set of orthogonal eigenvectors of \mathbf{L} associated with eigenvalue 0: that is, the set of indicator vectors $\{\mathbf{z}_\ell\}_{\ell=1, \dots, k}$ form a basis of $\text{span}(\mathbf{U}_k)$. Understanding arbitrary graphs with block structure as a perturbation of the ideal disconnected component case, the indicator vectors $\{\mathbf{z}_\ell\}_{\ell=1, \dots, k}$ of the blocks should live close to $\text{span}(\mathbf{U}_k)$ (in the sense that the difference between any \mathbf{z}_ℓ and its orthogonal projection onto $\text{span}(\mathbf{U}_k)$ is small). This in turn implies that every \mathbf{z}_ℓ should be approximately k -bandlimited.

As we will see next, the bandlimitedness assumption is very useful because it enables us to make use of generalized versions of Nyquist–Shannon sampling theorems, taking into account the graph.

5.5.3.2 Graph Sampling Theorems

The periodic sampling paradigm of the Shannon theorem for classical bandlimited signals does not apply to graphs without specific regular structure. In fact, a number of sampling schemes have been recently developed with the purpose of generalizing sampling theorems to graph signals [30, 109, 117, 134] (see [88] for a review of existing schemes).

Let us introduce some notations. Sampling entails selecting a set $\mathbf{S} = (s_1, \dots, s_m)$ of m nodes of the graph. To each possible sampling set, we associate a measurement matrix $\mathbf{M} = (\boldsymbol{\delta}_{s_1} | \boldsymbol{\delta}_{s_2} | \dots | \boldsymbol{\delta}_{s_m})^\top \in \mathbb{R}^{m \times n}$ where $\boldsymbol{\delta}_{s_i}(j) = 1$ if

$j = s_i$, and 0 otherwise. Now, consider a k -bandlimited signal $\mathbf{z} \in \text{span}(\mathbf{U}_k)$. The measurement of \mathbf{z} on \mathbf{S} reads:

$$\mathbf{y} = \mathbf{M}\mathbf{z} + \mathbf{n} \in \mathbb{R}^m, \quad (5.28)$$

where \mathbf{n} models measurement noise. The sampling question boils down to: how should we sample \mathbf{S} such that one can recover any bandlimited \mathbf{z} given its measurement \mathbf{y} ? There are three important components to this question: (i) how many samples m do we allow ourselves ($m = k$ being the strict theoretical minimum)? (ii) how much does it cost to sample? (iii) how do we in practice recover \mathbf{z} from \mathbf{y} and how much does that inversion cost?

There are a series of works that propose greedy algorithms to find the “best” set \mathbf{S} of minimal size $m = k$ that embed all k -bandlimited signals (see, for instance, [131] and references therein). These algorithms cost $\mathcal{O}(nk^4)$ and are thus not competitive in our setting.²² Moreover, in our case, we do not really need to be that strict on the number of samples and can allow more than k samples. A better choice is to use random graph sampling techniques. In the following we consider two types of independent sampling (uniform and leverage-score sampling) as well as a more involved method based on determinantal point processes.

Independent Sampling In the i.i.d. setting, one defines a discrete probability distribution $\mathbf{p} \in \mathbb{R}^n$ over the node set \mathbf{V} . The sampling set \mathbf{S} is then generated by drawing m nodes independently with replacement from \mathbf{p} . At each draw, the probability to sample node v_i is denoted by p_i . We have $\sum_i p_i = 1$ and write $\mathbf{P} = \text{diag}(p)$. Under this sampling scheme, the following restricted isometry property holds for the associated measurement matrix \mathbf{M} [109].

Theorem 5.9 *For any $\delta, \epsilon \in (0, 1)$, with probability at least $1 - \delta$:*

$$(1 - \epsilon)\|\mathbf{z}\|_2^2 \leq \frac{1}{m}\|\mathbf{M}\mathbf{P}^{-1/2}\mathbf{z}\|_2^2 \leq (1 + \epsilon)\|\mathbf{z}\|_2^2 \quad (5.29)$$

for all $\mathbf{z} \in \text{span}(\mathbf{U}_k)$ provided that

$$m \geq \frac{3}{\epsilon^2}(v_{\mathbf{p}}^k)^2 \log \frac{2k}{\delta} \quad (5.30)$$

where $v_{\mathbf{p}}^k$ is the so-called graph weighted coherence:

$$v_{\mathbf{p}}^k = \max_i \left\{ p_i^{-1/2} \|\mathbf{U}_k^{\top} \delta_i\|_2 \right\}. \quad (5.31)$$

This property is important as it says, in a nutshell, that any two different bandlimited signals will be identifiable post-sampling provided the number of samples is large

²²It takes longer to find a good sample than to run k -means on the whole dataset!

enough. The concept of large enough depends on $(v_{\mathbf{p}}^k)^2$: a measure of the interplay between the probability distribution and the norms of the rows of \mathbf{U}_k . In the uniform i.i.d. case since $p_i = 1/n$, one has $(v_{\mathbf{p}}^k)^2 = n \max_i \|\mathbf{U}_k^\top \delta_i\|_2^2$, which stays under control only for very regular graphs, but can be close to n in irregular graphs such as the star graph. The good news is that there exists an optimal sampling distribution (in the sense that it minimizes the right-hand side of inequality (5.30)) that adapts to the graph at hand:

$$p_i^* = \frac{\|\mathbf{U}_k^\top \delta_i\|_2^2}{k} \quad (5.32)$$

In fact, in this case, $(v_{\mathbf{p}^*}^k)^2$ matches its lower bound k and the necessary number of samples m to embed all bandlimited signals drops to $\mathcal{O}(k \log k)$. The distribution \mathbf{p}^* is also referred to by the name ‘‘leverage scores’’ in parts of the literature (see discussion in Sect. 5.3.1) [41]. As such, i.i.d. sampling under \mathbf{p}^* will be referred to as leverage score sampling.

Now, for lifting, there are several options.

- If one uses the unbiased decoder

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{w} \in \text{span}(\mathbf{U}_k)} \|\mathbf{P}_S^{-1/2}(\mathbf{M}\mathbf{w} - \mathbf{y})\|_2^2 \quad (5.33)$$

where $\mathbf{P}_S^{-1/2} = \mathbf{M}\mathbf{P}^{-1/2}\mathbf{M}^\top$, then the following reconstruction result holds [109]:

Theorem 5.10 *Let \mathbf{S} be the i.i.d. nodes sampled with distribution \mathbf{p} and \mathbf{M} be the associated sampling matrix. Let $\epsilon, \delta \in (0, 1)$ and suppose that m satisfies (5.30). With probability at least $1 - \delta$, for all $\mathbf{z} \in \text{span}(\mathbf{U}_k)$ and $\mathbf{n} \in \mathbb{R}^m$, the solution $\hat{\mathbf{z}}$ of (5.33) verifies:*

$$\|\hat{\mathbf{z}} - \mathbf{z}\|_2 \leq \frac{2}{\sqrt{m(1 - \epsilon)}} \|\mathbf{P}_S^{-1/2} \mathbf{n}\|_2.$$

This means that a noiseless measurement of a k -bandlimited signal yields a perfect reconstruction. Also, this quantifies how increasing m reduces the error of reconstruction due to a noisy measurement. Note that this error may be large if there is a significant measurement noise on a node that has a low probability of being sampled. However, by definition, this is not likely to happen.

- One can also use a label-propagation decoder reminiscent to semi-supervised learning techniques [15, 28]:

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{w} \in \mathbb{R}^n} \|\mathbf{P}_S^{-1/2}(\mathbf{M}\mathbf{w} - \mathbf{y})\|_2^2 + \gamma \mathbf{w}^\top g(\mathbf{L})\mathbf{w}, \quad (5.34)$$

where γ is a regularization parameter, $g(\mathbf{L})$ a graph filter operator as in (5.27) with $g(\lambda)$ a non-decreasing function. As g is non-decreasing, the regularization

term of (5.34) penalizes high frequency solutions, that is, solutions that are not smooth along paths of the graph. Theorems controlling the error of reconstruction are more involved and we refer the reader to Section 3.3 of [109] for details.

- Other decoders [14, 106] are in principle possible, replacing, for instance, the ℓ_2 Laplacian-based regularization $\mathbf{w}^\top g(\mathbf{L})\mathbf{w}$ by ℓ_1 -regularizers $\|\nabla\mathbf{w}\|_1$, but they come with an increased computation cost, lesser guarantees, and have not been used for spectral clustering: we will thus not detail them further.

Let us discuss the computation costs of the previous sampling and lifting techniques. In terms of sampling time, uniform sampling is obviously the most efficient and runs in $\mathcal{O}(k)$. Leverage score sampling is dominated by the computation of the optimal sampling distribution \mathbf{p}^* of (5.32), which takes $\mathcal{O}(nk)$ time.²³ In terms of lifting time, solving the decoder of (5.33) costs $\mathcal{O}(nk + mk^2)$. Solving the decoder of (5.34) costs $\mathcal{O}(et)$ via the conjugate gradient method, where t is the iteration number of the gradient solver (usually around 10 or 20 iterations suffice to obtain good accuracy when $g(\mathbf{L}) = \mathbf{L}$).

This discussion calls for a few remarks. First of all, these theorems are valid if we suppose that \mathbf{z} is exactly k -bandlimited, which is in fact only an approximation if we consider \mathbf{z} to be the ground truth indicator vectors of the k clusters to detect in the spectral clustering context. In this case, we can always decompose \mathbf{z} as the sum of its orthogonal projection onto $\text{span}(\mathbf{U}_k)$ and its complement $\boldsymbol{\beta}$: $\mathbf{z} = \mathbf{U}_k\mathbf{U}_k^\top\mathbf{z} + \boldsymbol{\beta}$. (5.28) becomes $\mathbf{y} = \mathbf{M}\mathbf{U}_k\mathbf{U}_k^\top\mathbf{z} + \mathbf{n}$ where \mathbf{n} now represents the sum of a measurement noise and the distance-to-model term $\mathbf{M}\boldsymbol{\beta}$. The aforementioned theorems can then be applied to $\mathbf{U}_k\mathbf{U}_k^\top\mathbf{z}$. Moreover, note that the decoder of (5.34) is not only faster than the other ones in general, it also does not constrain the solution $\hat{\mathbf{z}}$ to be exactly in $\text{span}(\mathbf{U}_k)$, which is in fact desirable in the spectral clustering context: we thus advocate for the decoder of (5.34).

DPP Sampling Determinantal point processes are a class of correlated random sampling strategies that strive to increase “diversity” in the samples, based on a kernel \mathbf{K} expliciting the similarity between variables. DPP sampling has been used successfully in a number of applications in machine learning (see, for instance, [74]).

Denote by $[n]$ the set of all subsets of $\{1, 2, \dots, n\}$. An element of $[n]$ could be the empty set, all elements of $\{1, 2, \dots, n\}$ or anything in between. DPPs are defined as follows:

Definition 5.4 (Determinantal Point Process [74]) Consider a point process, i.e., a process that randomly draws an element $\mathbf{S} \in [n]$. It is determinantal if, $\forall \mathbf{A} \subseteq \mathbf{S}$,

²³Note that the complexity is different from the leverage score computation of the Nyström techniques of Sects. 5.3.1 and 5.4.1 because, here, we suppose \mathbf{U}_k known whereas \mathbf{U}_k was not known in the previous sections. With \mathbf{U}_k known, computing the leverage scores only entails computing the normalized energy of each line of \mathbf{U}_k .

$$\mathbb{P}(\mathbf{A} \subseteq \mathbf{S}) = \det(\mathbf{K}_{\mathbf{A}}),$$

where $\mathbf{K} \in \mathbb{R}^{n \times n}$, a semi-definite positive matrix $0 \preceq \mathbf{K} \preceq \mathbf{1}$, is called the marginal kernel; and $\mathbf{K}_{\mathbf{A}}$ is the restriction of \mathbf{K} to the rows and columns indexed by the elements of \mathbf{A} .

The marginal probability π_i of sampling an element i is thus \mathbf{K}_{ii} . Consider the following projective kernel:

$$\mathbf{K}_k = \mathbf{U}_k \mathbf{U}_k^{\top}. \quad (5.35)$$

One can show that DPP samples from such projective kernels are necessarily of size k . After measuring the k -bandlimited signal \mathbf{z} on a DPP sample \mathbf{S} , one has the choice between the same decoders as before (see Eqs. (5.33) and (5.34)). For instance:

Theorem 5.11 *For all $\mathbf{z} \in \text{span}(\mathbf{U}_k)$, let $\mathbf{y} = \mathbf{M}\mathbf{z} + \mathbf{n} \in \mathbb{R}^k$ be a noisy measurement of \mathbf{z} on a DPP sample obtained from kernel \mathbf{K}_k . The decoder of (5.33) with $\mathbf{P} = \text{diag}(\pi_1, \dots, \pi_n)$ necessarily enables perfect reconstruction up to the noise level. Indeed, one obtains:*

$$\|\hat{\mathbf{z}} - \mathbf{z}\|_2 \leq \frac{1}{\sqrt{\lambda_{\min}(\mathbf{U}_k^{\top} \mathbf{M}^{\top} \mathbf{P}_{\mathbf{S}}^{-1} \mathbf{M} \mathbf{U}_k)}} \|\mathbf{P}_{\mathbf{S}}^{-1/2} \mathbf{n}\|_2. \quad (5.36)$$

Proof The proof is only partly in [131] and we complete it here. Let us write $\mathbf{z} = \mathbf{U}_k \alpha$. Solving (5.33) entails computing $\hat{\alpha} \in \mathbb{R}^k$ s.t. $\|\mathbf{P}_{\mathbf{S}}^{-1/2} (\mathbf{M} \mathbf{U}_k \hat{\alpha} - \mathbf{y})\|_2^2$ is minimal. Setting the derivative w.r.t. $\hat{\alpha}$ to 0, and replacing \mathbf{y} by $\mathbf{M} \mathbf{U}_k \alpha + \mathbf{n}$, yields:

$$\mathbf{U}_k^{\top} \mathbf{M}^{\top} \mathbf{P}_{\mathbf{S}}^{-1} \mathbf{M} \mathbf{U}_k \hat{\alpha} = \mathbf{U}_k^{\top} \mathbf{M}^{\top} \mathbf{P}_{\mathbf{S}}^{-1} \mathbf{M} \mathbf{U}_k \alpha + \mathbf{U}_k^{\top} \mathbf{M}^{\top} \mathbf{P}_{\mathbf{S}}^{-1} \mathbf{n}.$$

Recall that \mathbf{S} is a sample from a DPP with kernel \mathbf{K}_k : $\det(\mathbf{M} \mathbf{U}_k \mathbf{U}_k^{\top} \mathbf{M}^{\top})$ is thus strictly superior to 0, which implies that $\mathbf{M} \mathbf{U}_k$ is invertible, which in turn implies that $\hat{\alpha} = \alpha + (\mathbf{M} \mathbf{U}_k)^{-1} \mathbf{n}$. One thus has $\|\hat{\mathbf{z}} - \mathbf{z}\|_2 = \|\hat{\alpha} - \alpha\|_2 = \|(\mathbf{M} \mathbf{U}_k)^{-1} \mathbf{n}\|_2 = \|(\mathbf{P}_{\mathbf{S}}^{-1/2} \mathbf{M} \mathbf{U}_k)^{-1} \mathbf{P}_{\mathbf{S}}^{-1/2} \mathbf{n}\|_2$. Using the matrix 2-norm to bound this error yields

$$\|\hat{\mathbf{z}} - \mathbf{z}\|_2 \leq \sqrt{\lambda_{\max} \left[\left(\mathbf{U}_k^{\top} \mathbf{M}^{\top} \mathbf{P}_{\mathbf{S}}^{-1} \mathbf{M} \mathbf{U}_k \right)^{-1} \right]} \|\mathbf{P}_{\mathbf{S}}^{-1/2} \mathbf{n}\|_2,$$

as claimed.

Several comments are in order:

- The particular choice of kernel $\mathbf{K}_k = \mathbf{U}_k \mathbf{U}_k^{\top}$ implies that the marginal probability of sampling node v_i , $\pi_i = \|\mathbf{U}_k^{\top} \delta_i\|_2^2$, is proportional to the leverage scores p_i^* . The major difference between the i.i.d. leverage score approach and the DPP approach

comes from the negative correlations induced by the DPP. In fact, the probability of jointly sampling nodes v_i and v_j in the DPP case is $\pi_i\pi_j - \mathbf{K}_{ij}^2 = \pi_i\pi_j - (\delta_i^\top \mathbf{U}_k \mathbf{U}_k^\top \delta_j)^2$. The interaction term $(\delta_i^\top \mathbf{U}_k \mathbf{U}_k^\top \delta_j)^2$ will be typically large if v_i and v_j are in the same cluster, and small if not. In other words, different from the i.i.d. leverage score case where each new sample is drawn regardless of the past, the DPP procedure avoids to sample nodes containing redundant information.

- Whereas the leverage score approach only guarantees a RIP with high probability after $\mathcal{O}(k \log k)$ samples, the DPP approach has a stronger deterministic guarantee: it enables perfect invertibility (up to the noise level) after precisely $m = k$ samples. The reconstruction guarantee of (5.36) is nevertheless not satisfactory: even corrected by the marginal probabilities \mathbf{P}_S , the matrix $\mathbf{U}_k^\top \mathbf{M}^\top \mathbf{P}_S^{-1} \mathbf{M} \mathbf{U}_k$ can still have a very small λ_{\min} , such that reconstruction may be quite sensitive to noise. Improving this control is still an open problem. In practice, sampling independently 2 or 3 times from a DPP with kernel \mathbf{K}_k creates a set S of size $2k$ or $3k$ that is naturally more robust to noise.
- Whereas independent sampling is straightforward, sampling from a DPP with arbitrary kernel costs in general $\mathcal{O}(n^3)$ (see Algorithm 1 of [74] due to [62]). Thankfully, in the case of a projective kernel such as \mathbf{K}_k , one can sample a set in $\mathcal{O}(nk^2)$ based on Algorithm 3 of [132].

5.6 Perspectives

Almost two decades have passed since spectral clustering was first introduced. Since then, a large body of work has attempted to accelerate its computation. *So, has the problem been satisfactorily addressed?*—or, despite all these works, is there still room for improvement and further research?

To answer, we must first define what “satisfactorily addressed” would entail. As we have seen, the prototypical spectral clustering algorithm can be divided into three sub-problems: the similarity graph computation runs in $\mathcal{O}(dn^2)$; the spectral embedding computation runs in $\mathcal{O}(t(ek + nk^2))$ using an Arnoldi algorithm with t implicit restarts and assuming that e is the number of edges; and the k -means step runs in $\mathcal{O}(tnk^2)$, with t now being a bound on the number of iterations of the Lloyd-Max algorithm. Our criteria for evaluating an approximation algorithm aiming to accelerate one (or more) of these sub-problems are two-fold:

- We ask that the approximation algorithm’s computation cost is effectively lighter than the cost of the sub-problem(s) it is supposed to accelerate! The ultimate achievement is an order-of-magnitude improvement w.r.t. n (or e), d and/or k , especially when the complexity has no hidden constants (i.e., the algorithm is practically implementable). When such a gain is not possible, a gain on the constants of the theoretical cost is also considered worthwhile.
- The algorithm should come with convincing guarantees in terms of the quality of the found solution. Heuristics or partially motivated methods do not cut it. We

require that, under *mild assumptions*, the proposed solution is *provably close to* the exact solution. Let us clarify two aspects of this statement further:

- It is difficult to concretely classify assumptions as mild, but a useful rule of thumb is checking whether the theoretical results are meaningful for the significant majority of cases where spectral clustering would be used.
- The control of the approximation error comes in different flavors, that we detail here from the tightest to the loosest. The best possible error control in our context is a control over the clustering solution itself, via error measures such as the misclustering rate. This is unfortunately unrealistic in many cases. An excellent alternative is the multiplicative error—considered as the gold standard in approximation theory—over the k -means cost,²⁴ ensuring that the cost of the approximation is not larger than $1 + \epsilon$ times the cost of the exact solution. Next comes the additive error over the cost: ensuring that the cost difference between approximated and exact solutions is not larger than ϵ . All these error controls are referred to as end-to-end controls, and represent the limit of what we will consider a *satisfactory* error control.

Reviewing the literature, we were surprised to discover that there are rarely any algorithms meeting fully the proposed criteria: a faster algorithm with end-to-end control over the approximation error under mild assumptions. Let us revisit one by one the different approaches presented in Sects. 5.3, 5.4, and 5.5 examining them in light of our criteria for success. In each category of approximation algorithms, we order the methods according to the power of their error control.

Sampling Methods in the Original Feature Space [Sect. 5.3]

- *Representative points methods* as described in [63, 147] allow for an end-to-end control on the miss-clustering rate ρ , which is unfortunately quite loose. The constants involved in Theorem 5.3 are in fact undefined—thus potentially large—which is problematic knowing that ρ is by definition between 0 and 1. Also, the theorem’s assumptions include independence of the ϵ_i , which is hard to justify in practice. On the other hand, the computation gain of such methods is very appealing.
- *Feature projection methods*, where the dimension d of the original feature space is reduced to a dimension $d' \leq d$ based on Johnson–Lindenstrauss arguments, come with a multiplicative error control on the pairwise distances in the original feature space, thus providing a control on the obtained kernel matrix. The impact of this initial approximation on the final clustering result has not been studied.
- *Nyström-inspired methods* [19, 48, 85, 97] can be very efficient in practice especially because they do not need to build the graph. However, precisely because they do not build the graph, these methods cannot exactly perform two key

²⁴A control in terms of the k -means cost is usually considered as k -means is the last step of spectral clustering. Nevertheless, recalling the minimum cut perspective of Sect. 5.2.2, the control should arguably be in terms of rcut or ncut costs.

parts of the prototypical spectral clustering algorithm: the k -NN sparsification and the exact degree computation. The partial knowledge and sparsity of the kernel matrix also makes sampling difficult, as using leverage scores sampling is not possible anymore, whereas most other sampling schemes do not work very well with sparse matrices and come with weak guarantees. To the extent of our knowledge, there is also no convincing mathematical argument proving that using these methods will yield a clustering that is of similar quality to that produced by the exact spectral clustering algorithm.

- *Sketching methods* such as the random Fourier features [110] is yet another way of obtaining a pointwise multiplicative $(1 + \epsilon)$ error on the Gaussian kernel computation. RFF enable to compute a provably good low-rank approximation of the kernel. They nevertheless suffer from the same problems as Nyström-based techniques: without building the graph, sparsification and degree-normalization are uncontrolled. In addition, the guarantees on the low-rank approximation of the kernel do not transfer easily to guarantees of approximation of the spectral embedding \mathbf{U}_k .
- *Approximate nearest neighbors methods* are numerous and varied, and come with different levels of guarantees. Practical implementations of algorithms, however, often set aside theoretical guarantees to gain on efficiency and performances; and comparisons are usually done on benchmarks rather than on theoretical performances. In the best of cases, there is a control on how close the obtained nearest neighbor similarity graph is to the exact one, but with no end-to-end control.

Spectral Embedding Approximation Methods [Sect. 5.4]

- *Random eigenspace projection* is a very fast method and has been rigorously analyzed [95, 105, 111, 130]. It is true that a successful application depends on obtaining a good estimate of the k -th eigenvalue, which is very hard when the k -th eigenvalue gap is relatively small. Nevertheless, our current understanding of spectral clustering suggests that it only works well when the gap is (at least) moderately large. As such, though there are definitely situations in which random eigenspace projection will fail to provide an acceleration, these correspond to cases where one should not be using spectral clustering in the first place. The same argumentation can also be used in defense of all methods that come with mild gap assumptions (see coarsening and spectral sparsification).
- *Simple coarsening methods*, such as the heavy-edge matching heuristic [69], have nearly linear complexity, seem to work well in practice, and are accompanied by end-to-end additive error control [91]. Nevertheless, the current analysis of these heuristics only accounts for very moderate reductions ($m \geq n/2$) and thus does not fully prove their success: in real implementations coarsening is used in a multi-level fashion resulting to a drastic decrease in the graph size ($m = \mathcal{O}(n/2^c)$ for c levels), whereas the end-to-end control only works for a single level.
- *Advanced coarsening methods*, such as local variation methods [90], come with much stronger guarantees that allow for drastic size reduction and acceleration. Yet, thus far, all evidence suggests that finding a good enough coarsening is

computationally as hard as solving the spectral clustering problem itself. As a consequence, it is at this point unclear whether these methods can be used to accelerate spectral clustering.

- *Spectral sparsification techniques* come with excellent guarantees in theory: one may prove that a spectral sparsifier can be computed in nearly linear time and, moreover, the latter's spectrum will be provably close to the original one. Yet, we have reasons to doubt their practicality. Indeed, current algorithms are very complex, feature impractically large constants, and are only relevant for dense graphs. In addition, spectral sparsifiers, by definition, approximate the entire spectrum of a graph Laplacian matrix. However, spectral clustering only needs an approximation of a tiny fraction of the spectrum. From that perspective, it is reasonable to conclude that without modification current approaches will not yield the best possible approximation.
- *Nyström-approximation* applied directly to the Laplacian matrix is a good option, especially when combined with leverage score sampling. Nevertheless, an end-to-end error control has only been partially derived and is not yet satisfactory.

Sampling to Accelerate the k -Means Step [Sect. 5.5]

- *Exact methods to accelerate the Lloyd-Max algorithm*, may they be via avoiding unnecessary distance calculations or via a careful initialization are always useful and should be taken into account.
- *Coresets* come with the strongest guarantees: the minimum number of samples to guarantee a $(1 + \epsilon)$ multiplicative error on the cost function has been well studied. Nevertheless, practical coreset-sampling methods are scarce; and in the best cases, the sampling cost is of the same order of the Lloyd-Max running cost itself.
- *Graph-based sampling* comes with strong guarantees, but not over the k -means cost: on the reconstruction error based on a k -bandlimited model that is only an approximation in practice. Moreover, we interpret the reduced indicator vectors \mathbf{z}_ℓ^r obtained by running Lloyd-Max on the sampled set \mathbf{S} as (possibly noisy) measurements of \mathbf{z}_ℓ on \mathbf{S} . This interpretation currently lacks solid theoretical ground and impedes an end-to-end control of this approximation method. Nevertheless, the leverage-score-based sampling allows for a reduction in order of magnitude of the Lloyd-Max running cost.
- *Other methods* to accelerate k -means are not always appropriate to the spectral clustering context. Spectral feature dimension reduction is unnecessary in our context where $d = k$, sketching methods appropriate to distributed cases where n is very large are not appropriate neither as the spectral features need centralized data to be computed in any case.

In Practice The attentive reader will have remarked that, unsurprisingly, the tighter the error control, the more expensive the computation, and vice versa. Also, although we have put here an emphasis on the approximation error controls, it should not undermine the fact that methods from the whole spectrum are in practice useful, depending on the situation at hand, and specifically on the range of values of

n , d , and k . In very large d situations, a first step of random projection (or feature selection if some features are suspected to be too noisy) should be considered. Then, in situations where the exact computation of the proximity graph is too expensive, one may resort either to sketching methods or to Nyström-type methods to decrease the cost from quadratic to linear in n , and directly obtain an approximation of the spectral embedding without any explicit graph construction. These methods, however, do not take into account a sparsity constraint on the proximity graph and are usually rough on the degree correction they make.

The role of the sparsity constraint is not well understood theoretically, but seems to be important in some practical cases [138]. In such instances, a better option is to use approximate nearest neighbors methods to create a sparse similarity graph, and work from there. In extremely large data, say $n \geq 10^8$, the only workable methods are the representative-based, with, if possible, a first k -means (or compressive k -means [70]) to reduce n to m , or, in last resort, a uniform random sampling strategy.

In situations where one has to deal with such a large similarity graph that Arnoldi iterations are too expensive to compute the spectral embedding (either a graph created via approximate nearest neighbors or if the original data *is* a graph), projection methods such as in [21, 130], coarsening methods such as in [90], or Nyström-based methods are different possibilities.

Sampling methods to accelerate the last k -means step may seem to be a theoretical endeavor given that the Lloyd-Max algorithm is already very efficient. Due to the quadratic term in k , it is nevertheless in practice useful when k grows large. In this situation, hierarchical k -means [103] is a nice option. Coresets, because they are so stringent on the error control, have a hard time actually accelerating k -means, unless hybrid coreset-inspired methods are envisioned [49]. Finally, graph-based methods, because they take into account that spectral features are in fact derived from the graph itself, enable significant acceleration and are well-suited to the spectral clustering context.

Future Research Different directions of research could be envisioned to improve the state of the art:

- For Nyström-inspired methods in the context of Sect. 5.3 (directly applied on the original data) as well as the other methods based on computing a low-rank approximation of the kernel matrix \mathbf{K} , further work is needed to control both the sparsification and the degree correction, in order to bridge the gap between a provably good low-rank approximation of \mathbf{K} to a provably good low-rank approximation of \mathbf{R} .
- For Nyström methods in the context of Sect. 5.4 (applied on a known or well-approximated similarity graph), it would be interesting to extend Theorem 5.2 (for instance) to a control over $\|\mathbf{A}_k - \tilde{\mathbf{A}}_k\|$ instead of $\|\mathbf{A} - \tilde{\mathbf{A}}\|$. This would enable a tighter use of Davis–Kahan’s perturbation theorem in the discussion of Sect. 5.4.1 and, *in fine*, a better end-to-end guarantee.
- Projection-based methods of Sect. 5.4.3.2 currently necessitate to compute a value λ_* known to be in the interval $[\lambda_k, \lambda_{k+1})$. The algorithm used to do so is based on eigencount techniques that turn out to require as much computation time

as the Lanczos iterations needed to compute U_k exactly. One should relax this constraint to obtain end-to-end guarantees as a function of the distance between a coarsely estimated λ_* and the target interval.

- The derivation and analysis of randomized multi-level coarsening schemes with end-to-end guarantees is very much an open problem. We suspect that by utilizing spectrum-dependent sampling-schemes akin to leverage-scores one should be able to achieve results superior to heavy-edge matching in nearly linear time.
- There is an interesting similarity between coresets techniques and the graph-based sampling strategies discussed in Sect. 5.5.3 and it would be interesting to investigate this link theoretically, maybe paving the way to coresets for spectral clustering?

Finally, accelerating the prototypical spectral algorithm depicted in Algorithm 1 should not be the sole objective of researchers in this field. Indeed, taking the graph cut point-of-view of Sect. 5.2.2, Algorithm 1 makes three insufficiently motivated choices: (i) To begin with, the sparsification step in Algorithm 1 is not well understood. Apart from the fact that it is always computationally more convenient to work with a sparse similarity graph than a dense one, the precise effect of sparsification on the clustering performance has not been analyzed. (ii) As mentioned in Sect. 5.2.2.3, the relaxation employed by spectral relaxation is not unique. Why should we focus our attention on this one versus another? See, for instance, [24, 112] for recent alternative options. (iii) Finally, the use of k -means on the spectral features is not yet fully justified. Most of the end-to-end guarantees presented here compare the k -means cost of the exact solution to the k -means cost of the approximate solution. Given that the very use of k -means is not theoretically grounded, this choice of guarantee is debatable. Other options such as a control over the rcut or ncut objectives are possible (as in [96]) and should be further investigated.

Acknowledgements This work was partially supported by the CNRS PEPS I3A (Project RW4SPEc) and the Swiss National Science Foundation (Project DLEG, grant PZ00P2 179981).

References

1. Achlioptas, D.: Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. Syst. Sci.* **66**(4), 671–687 (2003)
2. Ali, H.T., Couillet, R.: Improved spectral community detection in large heterogeneous networks. *J. Mach. Learn. Res.* **18**(225), 1–49 (2018)
3. Altschuler, J., Bhaskara, A., Fu, G., Mirrokni, V., Rostamizadeh, A., Zadimoghaddam, M.: Greedy column subset selection: new bounds and distributed algorithms. In: Balcan, M.F., Weinberger, K.Q. (eds.) *Proceedings of the 33rd International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 48, pp. 2539–2548, New York (2016)

4. Anagnostopoulos, E., Emiris, I.Z., Psarros, I.: Low-quality dimension reduction and high-dimensional approximate nearest neighbor. In: Arge, L., Pach, J. (eds.) 31st International Symposium on Computational Geometry (SoCG 2015). Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, vol. 34, pp. 436–450 (2015). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik
5. Andoni, A., Indyk, P.: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In: 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06) (2006)
6. Arthur, D., Vassilvitskii, S.: k-means++: the advantages of careful seeding. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1027–1035. Society for Industrial and Applied Mathematics, Philadelphia (2007)
7. Arya, S., Mount, D.M., Netanyahu, D.M., Silverman, R., Wu, A.Y.: An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM* **45**(6), 891–923 (1998)
8. Aumüller, M., Bernhardsson, E., Faithfull, A.: ANN-benchmarks: a benchmarking tool for approximate nearest neighbor algorithms. In: Beecks, C., Borutta, F., Kröger, P., Seidl, T. (eds.) Similarity Search and Applications, Cham, pp. 34–49, Springer, Berlin (2017)
9. Avrithis, Y., Emiris, I.Z., Samaras, I.Z.: High-dimensional approximate nearest neighbor: k-d generalized randomized forests. arXiv:1603.09596 [cs] (2016)
10. Bachem, O., Lucic, M., Hassani, H., Krause, A.: Fast and provably good seedings for k-means. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 29, pp. 55–63. Curran Associates, New York (2016)
11. Bachem, O., Lucic, M., Krause, A.: Practical coresets constructions for machine learning. arXiv:1703.06476 [stat] (2017)
12. Bahmani, B., Moseley, B., Vattani, A., Kumar, R., Vassilvitskii, S.: Scalable K-means++. *Proc. VLDB Endow.* **5**(7), 622–633 (2012)
13. Bai, Z., Demmel, J., Dongarra, J., Ruhe, A., van der Vorst, H.: Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide. SIAM, Philadelphia (2000)
14. Bellec, P.C., Salmon, J., Vaiter, S.: A sharp oracle inequality for graph-slope. *Electron. J. Stat.* **11**(2), 4851–4870 (2017)
15. Bengio, Y., Delalleau, O., Roux, N.L.: Label propagation and quadratic criterion. In: Semi-Supervised Learning, pp. 193–216. MIT Press, Cambridge (2006)
16. Bhatia, R.: Matrix Analysis, vol. 169. Springer Science & Business Media, New York (1997)
17. Bie, T.D., Cristianini, N.: Fast SDP relaxations of graph cut clustering, transduction, and other combinatorial problems. *J. Mach. Learn. Res.* **7**, 1409–1436 (2006)
18. Bordenave, C., Lelarge, M., Massoulié, L.: Non-backtracking spectrum of random graphs: community detection and non-regular ramanujan graphs. In: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, pp. 1347–1357 (2015)
19. Bouneffouf, D., Birol, I.: Sampling with minimum sum of squared similarities for Nystrom-based large scale spectral clustering. In: The International Joint Conference on Artificial Intelligence (IJCAI 2015), pp. 2313–2319 (2015)
20. Boutsidis, C., Drineas, P., Mahoney, M.W.: Unsupervised feature selection for the k-means clustering problem. In: Bengio, Y., Schuurmans, D., Lafferty, J.D., Williams, C.K.I., Culotta, A. (eds.) Advances in Neural Information Processing Systems, vol. 22, pp. 153–161. Curran Associates, New York (2009)
21. Boutsidis, C., Gittens, A., Kambadur, P.: Spectral clustering via the power method-provably. In: International Conference on Machine Learning (ICML) (2015)
22. Boutsidis, C., Zouzias, A., Mahoney, M.W., Drineas, P.: Randomized dimensionality reduction for k -means clustering. *IEEE Trans. Inf. Theory* **61**(2), 1045–1062 (2015)
23. Brand, M., Huang, K.: A unifying theorem for spectral embedding and clustering. In: 9th International Conference on Artificial Intelligence and Statistics, AISTATS (2003)

24. Bresson, X., Laurent, T., Uminsky, D., von Brecht, J.: Multiclass total variation clustering. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 26, pp. 1421–1429. Curran Associates, New York (2013)
25. Cai, D., Zhang, C., He, X.: Unsupervised feature selection for multi-cluster data. In: *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD'10*, Washington, p. 333. ACM Press, New York (2010)
26. Calvetti, D., Reichel, L., Sorensen, D.C.: An implicitly restarted Lanczos method for large symmetric eigenvalue problems. *Electron. Trans. Numer. Anal.* **2**(1), 21 (1994)
27. Celebi, M.E., Kingravi, H.A., Vela, P.A.: A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Syst. Appl.* **40**(1), 200–210 (2013)
28. Chapelle, O., Schlkopf, B., Zien, A.: *Semi-Supervised Learning*, 1st edn. MIT Press, Cambridge (2010)
29. Chen, X., Cai, D.: Large scale spectral clustering with landmark-based representation. In: *AAAI Conference on Artificial Intelligence* (2011)
30. Chen, S., Varma, R., Sandryhaila, A., Kovacevic, J.: Discrete signal processing on graphs: sampling theory. *CoRR*, abs/1503.05432 (2015)
31. Chitta, R., Jin, R., Jain, A.K.: Efficient kernel clustering using random Fourier features. In: *2012 IEEE 12th International Conference on Data Mining*, pp. 161–170 (2012)
32. Chung, F.R., Graham, F.C.: *Spectral Graph Theory*, vol. 92. American Mathematical Society, Providence (1997)
33. Cohen, M.B., Elder, S., Musco, C., Musco, C., Persu, M.: Dimensionality reduction for k-means clustering and low rank approximation. In: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC'15*, New York, pp. 163–172. ACM, New York (2015)
34. Dall'Amico, L., Couillet, R., Tremblay, N.: Revisiting the Bethe-Hessian: improved community detection in sparse heterogeneous graphs. In: *NeurIPS* (2019)
35. Dash, M., Koot, P.W.: Feature selection for clustering. In: Liu, L., özsu, M.T. (eds.) *Encyclopedia of Database Systems*, pp. 1119–1125. Springer, Boston (2009)
36. Davis, C.: The rotation of eigenvectors by a perturbation. *J. Math. Anal. Appl.* **6**(2), 159–173 (1963)
37. Deshpande, A., Rademacher, L., Vempala, S., Wang, G.: Matrix approximation and projective clustering via volume sampling. In: *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, pp. 1117–1126. Society for Industrial and Applied Mathematics, Philadelphia (2006)
38. Dhillon, I.S., Guan, Y., Kulis, B.: Weighted graph cuts without eigenvectors a multilevel approach. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(11), 1944–1957 (2007)
39. Di Napoli, E., Polizzi, E., Saad, Y.: Efficient estimation of eigenvalue counts in an interval. *Numer. Linear Algebra Appl.* **23**(4), 674–692 (2016)
40. Dong, W., Moses, C., Li, K.: Efficient k-nearest neighbor graph construction for generic similarity measures. In: *Proceedings of the 20th International Conference on World Wide Web - WWW'11*, Hyderabad, p. 577. ACM Press, New York (2011)
41. Drineas, P., Mahoney, M.W.: *Lectures on randomized numerical linear algebra*. *Math. Data* **25**, 1 (2018)
42. Drineas, P., Frieze, A.M., Kannan, R., Vempala, S., Vinay, V.: Clustering in large graphs and matrices. In: *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, vol. 99, pp. 291–299. Citeseer (1999)
43. Drineas, P., Kannan, R., Mahoney, M.W.: Fast Monte Carlo algorithms for matrices ii: computing a low-rank approximation to a matrix. *SIAM J. Comput.* **36**(1), 158–183 (2006)
44. Duan, R., Pettie, S.: Linear-time approximation for maximum weight matching. *J. ACM* **61**(1), 1 (2014)
45. Feldman, D., Langberg, M.: A unified framework for approximating and clustering data. In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, pp. 569–578. ACM, New York (2011)

46. Fiedler, M.: Algebraic connectivity of graphs. *Czechoslov. Math. J.* **23**(2), 298–305 (1973)
47. Fortunato, S.: Community detection in graphs. *Phys. Rep.* **486**(3–5), 75–174 (2010)
48. Fowlkes, C., Belongie, S., Chung, F., Malik, J.: Spectral grouping using the Nystrom method. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(2), 214–225 (2004)
49. Frahling, G., Sohler, C.: A fast k-means implementations using coresets. *Int. J. Comput. Geom. Appl.* **18**(6), 605–625 (2008)
50. Frieze, A., Kannan, R., Vempala, S.: Fast Monte-Carlo algorithms for finding low-rank approximations. *J. ACM* **51**(6), 1025–1041 (2004)
51. Fu, C., Xiang, C., Wang, C., Cai, D.: Fast approximate nearest neighbor search with the navigating spreading-out graph. arXiv:1707.00143 [cs] (2017)
52. Gadde, A., Anis, A., Ortega, A.: Active semi-supervised learning using sampling theory for graph signals. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD'14, New York*, pp. 492–501. ACM, New York (2014)
53. Gadde, A., Gad, E.E., Avestimehr, S., Ortega, A.: Active learning for community detection in stochastic block models. In: *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 1889–1893 (2016)
54. Gittens, A., Mahoney, M.W.: Revisiting the Nyström method for improved large-scale machine learning. *J. Mach. Learn. Res.* **17**(1), 3977–4041 (2016)
55. Gribonval, R., Blanchard, G., Keriven, N., Traonmilin, Y.: Compressive statistical learning with random feature moments. arXiv:1706.07180 [cs, math, stat] (2017)
56. Guattery, S., Miller, G.L.: On the performance of spectral graph partitioning methods. In: *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms SODA*, vol. 95, pp. 233–242 (1995)
57. Guattery, S., Miller, G.L.: On the quality of spectral separators. *SIAM J. Matrix Anal. Appl.* **19**(3), 701–719 (1998)
58. Halko, N., Martinsson, P., Tropp, J.: Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* **53**(2), 217–288 (2011)
59. Hamerly, G., Drake, J.: *Accelerating Lloyd's Algorithm for k-Means Clustering*, pp. 41–78. Springer, Cham (2015)
60. He, X., Cai, D., Niyogi, P.: Laplacian score for feature selection. In: *Advances in Neural Information Processing Systems*, pp. 507–514 (2006)
61. Hendrickson, B., Leland, R.W.: A multi-level algorithm for partitioning graphs. *Supercomputing* **95**(28), 1–14 (1995)
62. Hough, J.B., Krishnapur, M., Peres, Y., Virág, B.: Determinantal processes and independence. *Probab. Surv.* **3**, 206–229 (2006)
63. Huang, L., Yan, D., Taft, N., Jordan, M.I.: Spectral clustering with perturbed data. In: Koller, D., Schuurmans, D., Bengio, Y., Bottou, L. (eds.) *Advances in Neural Information Processing Systems*, vol. 21, pp. 705–712. Curran Associates, New York (2009)
64. Hunter, B., Strohmer, T., Simos, T.E., Psihoyios, G., Tsitouras, C.: Compressive spectral clustering, pp. 1720–1722, Rhodes (2010)
65. Isufi, E., Loukas, A., Simonetto, A., Leus, G.: Autoregressive moving average graph filtering. *IEEE Trans. Signal Process.* **65**(2), 274–288 (2017)
66. Kalantidis, Y., Avrithis, Y.: Locally optimized product quantization for approximate nearest neighbor search. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014)
67. Kanungo, T., Mount, D.M., Netanyahu, N.S., Piatko, C.D., Silverman, R., Wu, A.Y.: An efficient k-means clustering algorithm: analysis and implementation. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 881–892 (2002)
68. Karger, D.R.: Minimum cuts in near-linear time. *J. ACM* **47**(1), 46–76 (2000)
69. Karypis, G., Kumar, V.: A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.* **20**(1), 359–392 (1998)

70. Keriven, N., Tremblay, N., Traonmilin, Y., Gribonval, R.: Compressive K-means. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6369–6373 (2017)
71. Kolev, P., Mehlhorn, K.: A note on spectral clustering. arXiv preprint arXiv:1509.09188 (2015)
72. Koutis, I., Miller, G.L., Peng, R.: Approaching optimality for solving SDD linear systems. In: 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, pp. 235–244 (2010)
73. Krzakala, F., Moore, C., Mossel, E., Neeman, J., Sly, A., Zdeborová, L., Zhang, P.: Spectral redemption in clustering sparse networks. *Proc. Natl. Acad. Sci.* **110**(52), 20935–20940 (2013)
74. Kulesza, A., Taskar, B.: Determinantal point processes for machine learning. *Found. Trends Mach. Learn.* **5**(2–3), 123–286 (2012)
75. Kumar, A., Kannan, R.: Clustering with spectral norm and the k-means algorithm. In: 2010 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 299–308. IEEE, Piscataway (2010)
76. Kumar, A., Sabharwal, Y., Sen, S.: A simple linear time $(1+\epsilon)$ -approximation algorithm for k-means clustering in any dimensions. In: Proceedings of 45th Annual IEEE Symposium on Foundations of Computer Science, 2004, pp. 454–462. IEEE, Piscataway (2004)
77. Kumar, S., Mohri, M., Talwalkar, A.: Sampling techniques for the nystrom method. In: Artificial Intelligence and Statistics, pp. 304–311 (2009)
78. Kumar, S., Mohri, M., Talwalkar, A.: Sampling methods for the nyström method. *J. Mach. Learn. Res.* **13**, 981–1006 (2012)
79. Langberg, M., Schulman, L.J.: Universal ϵ -approximators for integrals. In: Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 598–607. SIAM, Philadelphia (2010)
80. Laurent, B., Massart, P.: Adaptive estimation of a quadratic functional by model selection. *Ann. Stat.* **28**, 1302–1338 (2000)
81. Le, Q., Sarlós, T., Smola, A.: Fastfood-approximating kernel expansions in loglinear time. In: Proceedings of the International Conference on Machine Learning, vol. 85 (2013)
82. Lee, H., Battle, A., Raina, R., Ng, A.Y.: Efficient sparse coding algorithms. In: Schölkopf, B., Platt, J.C., Hoffman, T. (eds.) *Advances in Neural Information Processing Systems*, vol. 19, pp. 801–808. MIT Press, Cambridge (2007)
83. Lee, J.R., Gharan, S.O., Trevisan, L.: Multiway spectral partitioning and higher-order cheeger inequalities. *J. ACM* **61**(6), 37 (2014)
84. Lei, J., Rinaldo, A.: Consistency of spectral clustering in stochastic block models. *Ann. Stat.* **43**(1), 215–237 (2015)
85. Li, M., Lian, X.-C., Kwok, J., Lu, B.-L.: Time and space efficient spectral clustering via column sampling. In: Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition, pp. 2297–2304. IEEE Computer Society, Washington (2011)
86. Li, Q., Liu, W., Li, L., Wang, R.: Towards large scale spectral problems via diffusion process. In: 2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA), pp. 1–7 (2017)
87. Lloyd, S.: Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **28**(2), 129–137 (1982)
88. Lorenzo, P., Barbarossa, S., Banelli, P.: Chapter 9 - sampling and recovery of graph signals. In M. Djurić, P., Richard, C. (eds.), *Cooperative and Graph Signal Processing*, pp. 261–282. Academic Press, Cambridge (2018)
89. Loukas, A.: How close are the eigenvectors of the sample and actual covariance matrices? In: Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, International Convention Centre, Sydney, 6–11, vol. 70, pp. 2228–2237 (2017)
90. Loukas, A.: Graph reduction with spectral and cut guarantees. arXiv preprint arXiv:1808.10650 (2018)

91. Loukas, A., Vndergheynst, P.: Spectrally approximating large graphs with smaller graphs. In: International Conference on Machine Learning (ICML) (2018)
92. Loukas, A., Simonetto, A., Leus, G.: Distributed autoregressive moving average graph filters. *IEEE Signal Process. Lett.* **22**(11), 1931–1935 (2015)
93. Makarychev, K., Makarychev, Y., Razenshteyn, I.: Performance of Johnson-Lindenstrauss transform for k-means and k-medians clustering. arXiv:1811.03195 [cs] (2018)
94. Malkov, Y.A., Yashunin, D.A.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. arXiv:1603.09320 [cs], p. 13 (2016)
95. Martin, L., Loukas, A., Vndergheynst, P.: Fast approximate spectral clustering for dynamic networks. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden. Proceedings of Machine Learning Research, vol. 80, pp. 3423–3432. PMLR (2018)
96. Mohan, M., Monteleoni, C.: Beyond the nystrom approximation: speeding up spectral clustering using uniform sampling and weighted kernel k-means. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence IJCAI (2017)
97. Mohan, M., Monteleoni, C.: Exploiting sparsity to improve the accuracy of nystrom-based large-scale spectral clustering. In: 2017 International Joint Conference on Neural Networks (IJCNN) (2017)
98. Muja, M., Lowe, D.G.: Scalable nearest neighbor algorithms for high dimensional data. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**, 2227–2240 (2014)
99. Munteanu, A., Schwiigelshohn, C.: Coresets-methods and history: a theoreticians design pattern for approximation and streaming algorithms. *Künstl. Intell.* **32**, 37–53 (2017)
100. Musco, C., Musco, C.: Recursive sampling for the nystrom method. In: Advances in Neural Information Processing Systems, pp. 3833–3845 (2017)
101. Newling, J., Fleuret, F.: Fast k-means with accurate bounds. In: International Conference on Machine Learning, pp. 936–944 (2016)
102. Ng, A., Jordan, M., Weiss, Y., et al.: On spectral clustering: analysis and an algorithm. *Adv. Neural Inf. Process. Syst.* **2**, 849–856 (2002)
103. Nister, D., Stewenius, H.: Scalable recognition with a vocabulary tree. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), vol. 2, pp. 2161–2168 (2006)
104. Ostrovsky, R., Rabani, Y., Schulman, L.J., Swamy, C.: The effectiveness of Lloyd-type methods for the k-means problem. In: 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), pp. 165–176 (2006)
105. Paratte, J., Martin, L.: Fast eigenspace approximation using random signals. arXiv preprint arXiv:1611.00938 (2016)
106. Pena, R., Bresson, X., Vndergheynst, P.: Source localization on graphs via ℓ_1 recovery and spectral graph theory. In: 2016 IEEE 12th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP), pp. 1–5 (2016)
107. Peng, R., Sun, H., Zanetti, L.: Partitioning well-clustered graphs: spectral clustering works! In: Conference on Learning Theory, pp. 1423–1455 (2015)
108. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8 (2007)
109. Puy, G., Tremblay, N., Gribonval, R., Vndergheynst, P.: Random sampling of bandlimited signals on graphs. *Appl. Comput. Harmonic Anal.* **44**(2), 446–475 (2018)
110. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: Advances in Neural Information Processing Systems, pp. 1177–1184 (2008)
111. Ramasamy, D., Madhoo, U.: Compressive spectral embedding: sidestepping the SVD. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 28, pp. 550–558. Curran Associates, New York (2015)

112. Rangapuram, S.S., Mudrakarta, P.K., Hein, M.: Tight continuous relaxation of the balanced k-cut problem. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 27, pp. 3131–3139. Curran Associates, New York (2014)
113. Ros, F., Guillaume, S.: ProTraS: A probabilistic traversing sampling algorithm. *Expert Syst. Appl.* **105**, 65–76 (2018)
114. Saade, A., Krzakala, F., Zdeborová, L.: Spectral clustering of graphs with the Bethe Hessian. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 27 pp. 406–414. Curran Associates, Inc. (2014)
115. Safro, I., Sanders, P., Schulz, C.: Advanced coarsening schemes for graph partitioning. *J. Exp. Algorithmics* **19**, 2 (2015)
116. Sakai, T., Imiya, A.: Fast spectral clustering with random projection and sampling. In: Perner, P. (ed) *Machine Learning and Data Mining in Pattern Recognition*. Lecture Notes in Computer Science, vol. 5632, pp. 372–384. Springer, Berlin (2009)
117. Sakiyama, A., Tanaka, Y., Tanaka, T., Ortega, T.: Eigendecomposition-free sampling set selection for graph signals. arXiv:1809.01827 [eess] (2018)
118. Sandryhaila, A., Moura, J.: Big data analysis with signal processing on graphs: representation and processing of massive data sets with irregular structure. *IEEE Signal Process. Mag.* **31**(5), 80–90 (2014)
119. Sculley, D.: Web-scale k-means clustering. In: *Proceedings of the 19th International Conference on World Wide Web - WWW'10*, Raleigh, p. 1177. ACM Press, New York (2010)
120. Sedley, D.: An introduction to Plato's theory of forms. *R. Inst. Philos. Suppl.* **78**, 3–22 (2016)
121. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(8), 888–905 (2000)
122. Shuman, D.I., Vandergheynst, P., Frossard, P.: Chebyshev polynomial approximation for distributed signal processing. In: *2011 International Conference on Distributed Computing in Sensor Systems and Workshops (DCOSS)*, pp. 1–8. IEEE, Piscataway (2011)
123. Shuman, D.I., Narang, S.K., Frossard, P., Ortega, A., Vandergheynst, P.: The emerging field of signal processing on graphs: extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Process. Mag.* **30**(3), 83–98 (2013)
124. Spielman, D.A., Srivastava, N.: Graph sparsification by effective resistances. *SIAM J. Comput.* **40**(6), 1913–1926 (2011)
125. Spielman, D.A., Teng, S.-H.: Spectral sparsification of graphs. *SIAM J. Comput.* **40**(4), 981–1025 (2011)
126. Stoer, M., Wagner, F.: A simple min-cut algorithm. *J. ACM* **44**(4), 585–591 (1997)
127. Sutherland, D.J., Schneider, J.: On the error of random Fourier features. arXiv:1506.02785 [cs, stat] (2015)
128. Tarsitano, A.: A computational study of several relocation methods for k-means algorithms. *Pattern Recognit.* **36**(12), 2955–2966 (2003)
129. Tremblay, N., Puy, G., Borgnat, P., Gribonval, R., Vandergheynst, P.: Accelerated spectral clustering using graph filtering of random signals. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2016)
130. Tremblay, N., Puy, G., Gribonval, R., Vandergheynst, P.: Compressive spectral clustering. In: *33rd International Conference on Machine Learning*, New York (2016)
131. Tremblay, N., Amblard, P.O., Barthelmé, S.: Graph sampling with determinantal processes. In: *2017 25th European Signal Processing Conference (EUSIPCO)*, pp. 1674–1678 (2017)
132. Tremblay, N., Barthelme, S., Amblard, P.-O.: Optimized algorithms to sample determinantal point processes. arXiv:1802.08471 [cs, stat] (2018)
133. Tremblay, N., Goncalves, P., Borgnat, P.: Chapter 11 - design of graph filters and filterbanks. In: Djurić, P.M., Richard, C. (eds.) *Cooperative and Graph Signal Processing*, pp. 299–324. Academic Press, Cambridge (2018)
134. Tsitsvero, M., Barbarossa, S., Di Lorenzo, P.: Signals on graphs: uncertainty principle and sampling. *IEEE Trans. Signal Process.* **64**(18), 4845–4860 (2016)

135. Ulrike von Luxburg, O.B., Belkin, M.: Consistency of spectral clustering. *Ann. Stat.* **36**(2), 555–586 (2008)
136. Vishnoi, N.K., et al.: $L_x = b$. *Found. Trends Theor. Comput. Sci.* **8**(1–2), 1–141 (2013)
137. Vladymyrov, M., Carreira-Perpinan, M.A.: Fast, accurate spectral clustering using locally linear landmarks. In: 2017 International Joint Conference on Neural Networks, pp. 3870–3879 (2017)
138. Von Luxburg, U.: A tutorial on spectral clustering. *Stat. Comput.* **17**(4), 395–416 (2007)
139. Wagner, D., Wagner, F.: Between min cut and graph bisection. In: Borzyszkowski, A.M., Sokolowski, S. (eds.), *Mathematical Foundations of Computer Science 1993*, Berlin, Heidelberg, pp. 744–750. Springer, Berlin (1993)
140. Wang, Y., Feng, Z.: Towards scalable spectral clustering via spectrum-preserving sparsification. *arXiv preprint arXiv:1710.04584* (2017)
141. Wang, L., Leckie, C., Ramamohanarao, K., Bezdek, J.: Approximate spectral clustering. In: Theeramunkong, T., Kijssirikul, B., Cercone, N., Ho, T.-B. (eds.), *Advances in Knowledge Discovery and Data Mining. Lecture Notes in Computer Science*, vol. 5476, pp. 134–146. Springer, Berlin (2009)
142. Wang, J., Wang, J., Ke, Q., Zeng, G., Li, S.: Fast approximate k -means via cluster closures. In: *Multimedia Data Mining and Analytics*, pp. 373–395. Springer, Berlin (2015)
143. Wei, Y.-C., Cheng, C.-K.: Towards efficient hierarchical designs by ratio cut partitioning. In: 1989 IEEE International Conference on Computer-Aided Design (ICCAD), pp. 298–301 (1989)
144. Wierzchoń, S.T., Kłopotek, M.A.: Spectral clustering. In: *Modern Algorithms of Cluster Analysis*, pp. 181–259. Springer, Cham (2018)
145. Woodruff, D.P., et al.: Sketching as a tool for numerical linear algebra. *Found. Trends Theor. Comput. Sci.* **10**(1–2), 1–157 (2014)
146. Wu, L., Chen, P.-Y., Yen, I. E.-H., Xu, F., Xia, Y., Aggarwal, C.: Scalable spectral clustering using random binning features. *arXiv:1805.11048 [cs, stat]* (2018)
147. Yan, D., Huang, D., Jordan, M.I.: Fast approximate spectral clustering. In: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD'09, Paris*, pp. 907–916. ACM, New York (2009)
148. Zhang, K., Tsang, I.W., Kwok, J.T.: Improved nyström low-rank approximation and error analysis. In: *Proceedings of the 25th International Conference on Machine Learning*, pp. 1232–1239. ACM, New York (2008)
149. Zhao, Z., Liu, H.: Spectral feature selection for supervised and unsupervised learning. In: *Proceedings of the 24th International Conference on Machine Learning - ICML'07, Corvallis*, pp. 1151–1157. ACM Press, New York (2007)