
Power-Law Graph Cuts

Xiangyang Zhou

Jiaxin Zhang

Brian Kulis

Department of Computer Science and Engineering
The Ohio State University

Abstract

Algorithms based on spectral graph cut objectives such as normalized cuts, ratio cuts and ratio association have become popular in recent years because they are widely applicable and simple to implement via standard eigenvector computations. Despite strong performance for a number of clustering tasks, spectral graph cut algorithms still suffer from several limitations: first, they require the number of clusters to be known in advance, but this information is often unknown *a priori*; second, they tend to produce clusters with uniform sizes. In some cases, the true clusters exhibit a known size distribution; in image segmentation, for instance, human-segmented images tend to yield segment sizes that follow a power-law distribution. In this paper, we propose a general framework of power-law graph cut algorithms that produce clusters whose sizes are power-law distributed, and also does not fix the number of clusters upfront. To achieve our goals, we treat the Pitman-Yor exchangeable partition probability function (EPPF) as a regularizer to graph cut objectives. Because the resulting objectives cannot be solved by relaxing via eigenvectors, we derive a simple iterative algorithm to locally optimize the objectives. Moreover, we show that our proposed algorithm can be viewed as performing MAP inference on a particular Pitman-Yor mixture model. Our experiments on various data sets show the effectiveness of our algorithms.

1 Introduction

As one of the most fundamental problems in machine learning, clustering has received a considerable amount of attention and has applications in data mining, computer vision, statistics, social sciences, and others. Spectral graph cut algorithms such as normalized cuts [1], ratio cut [2] and ratio association [1, 3] are one of the most studied and utilized classes of clustering methods. These algorithms aim to cluster data by first constructing a similarity graph based on the given data, then “cutting” the graph into groups of nodes according to a graph-theoretic objective. Normalized cuts has been widely used in the computer vision community for image segmentation [1] and other problems [4] while ratio cut has been applied in circuit layout [2]. Though these graph cut problems can be shown to be NP-hard, several effective algorithms have been proposed, including eigenvector-based approaches [1] as well as methods based on kernel k-means [3].

Despite the success of spectral graph cut algorithms, they do suffer from several important limitations. For one, they require the number of clusters to be known before running the algorithm, but in many applications the number of clusters is not known *a priori*. More importantly, many graph cut objectives, such as the normalized cut objective and the ratio cut objective, favor clusters of equal size or degree, which typically leads these algorithms to produce clusters with nearly uniform sizes. Consider image segmentation, the canonical application of normalized cuts. As shown in [5], human-segmented images yield segments that are far from uniform; in fact, they follow a *power-law distribution* in terms of their segment sizes. Power-law distributions arise frequently in a number of other clustering applications as well. For instance, because income follows a power-law distribution, attempting to cluster individuals into income brackets using census data would likely fail when applying standard clustering techniques. Other phenomena exhibiting power-law distributions include the populations of cities, the intensities of earthquakes, and the sizes of

Appearing in Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS) 2015, San Diego, CA, USA. JMLR: W&CP volume 38. Copyright 2015 by the authors.

power outages [6]. These applications—and the lack of existing graph clustering methods that specifically encourage power-law cluster size structure—motivate our work.

In this paper, we propose a general framework of power-law graph cut algorithms that encourages cluster sizes to be power-law distributed, and does not fix the number of clusters upfront. To achieve both goals, we borrow ideas from Bayesian nonparametrics [7], which provide a principled way to automatically infer both the parameters of a model as well as its complexity. We observe that the Pitman-Yor Chinese restaurant process [8], a Bayesian nonparametric prior that generalizes the Chinese restaurant process, yields clusters whose sizes follow a power-law distribution. We treat the Pitman-Yor exchangeable partition probability function (EPPF) [9] as a *regularizer* for graph cut objectives, so that the resulting objectives favor clusters that both have a small graph cut objective value as well as a power-law cluster size structure.

Algorithmically, incorporating the Pitman-Yor EPPF into existing cut formulations results in an optimization problem where standard spectral methods are no longer applicable. Inspired by the connection between spectral graph cut objectives and weighted kernel k-means [3], we derive a simple k-means-like iterative algorithm to optimize several power-law graph cut objectives. As with k-means, our proposed algorithm is guaranteed to converge to a local optima in a finite number of steps. We further demonstrate that our graph cut problem may be viewed precisely as a MAP problem for a particular Pitman-Yor Gaussian mixture model. Finally, to demonstrate the utility of our algorithm, we perform extensive experiments using power-law normalized cuts on synthetic datasets, real-world data with power-law structure, and image segmentation.

Related Work: Small-variance asymptotics have recently been extended to Bayesian nonparametric models to yield simple k-means-like algorithms [10, 11]; one of the applications of that line of work is a normalized cut algorithm that does not fix the number of clusters upfront [10]. However, that approach cannot be directly applied to Pitman-Yor process mixture models, as small-variance asymptotics on the Pitman-Yor process model fail to capture any power-law characteristics.

The most related work to ours is [12], an algorithm for scalable power-law clustering based on adapting k-means. Specifically, the authors propose a power-law data clustering algorithm based on modifying the Pitman-Yor process and performing a small-variance asymptotic analysis on the modified Piman-Yor pro-

cess. However, their objective function does not guarantee the generation of a power-law distributed cluster sizes and the optimal clustering solutions for their objective are often trivial. We will discuss this method further in Section 4.4 and Section 5.

The work of [13] is complementary to ours in that it deals with graphs with power-law degree distributions while our paper applies to graphs with power-law cluster size distributions. Finally, the work of [5] introduces a model for segmentation based on Pitman-Yor priors, but it is specific to the image domain whereas our method is a general graph clustering algorithm.

2 Background

We begin with a brief discussion about spectral graph cut algorithms and their connection to weighted kernel k-means.

2.1 Spectral graph cut algorithms

In the graph clustering setting, we are given an undirected weighted graph $G = (\mathcal{V}, \mathcal{E})$, in which $\mathcal{V} = \{v_1, \dots, v_n\}$ denotes vertices and \mathcal{E} denotes edges. The weight of an edge between two vertices represents their similarity. The corresponding adjacency matrix A is a $|\mathcal{V}|$ -by- $|\mathcal{V}|$ matrix whose entry A_{ij} represents the weight of the edge between v_i and v_j .

The idea behind graph cuts is to partition the graph into k disjoint clusters such that the edges within a cluster have high weight and the edges between clusters have low weight. Several different graph cut objectives have been proposed [1, 2, 14], among which normalized cuts [1] and ratio cut [2] are two of the most popular. Denote

$$\text{cut}(\mathcal{V}_1, \mathcal{V}_2) = \sum_{i \in \mathcal{V}_1, j \in \mathcal{V}_2} A_{ij},$$

i.e., the sum of the edge weights between \mathcal{V}_1 and \mathcal{V}_2 , and

$$\text{deg}(\mathcal{V}_1) = \text{cut}(\mathcal{V}_1, \mathcal{V}),$$

the sum of all edge weights between \mathcal{V}_1 and \mathcal{V} . Normalized cuts (sometimes called k -way normalized cuts) aims to minimize the cut relative to the degree of the cluster. The objective can be expressed as

$$\text{NCut}(G) = \min_{\mathcal{V}_1, \dots, \mathcal{V}_k} \sum_{i=1}^k \frac{\text{cut}(\mathcal{V}_i, \mathcal{V}/\mathcal{V}_i)}{\text{deg}(\mathcal{V}_i)}.$$

While this objective can be shown to be NP-complete, a relaxation of it can be globally optimized using spectral methods by computing the first k eigenvalues of the normalized Laplacian constructed from the adjacency matrix A [15].

The ratio cut objective differs from normalized cuts in that it seeks to minimize the cut between clusters and the remaining vertices. It is expressed as

$$\text{RCut}(G) = \min_{\mathcal{V}_1, \dots, \mathcal{V}_k} \sum_{i=1}^k \frac{\text{cut}(\mathcal{V}_i, \mathcal{V}/\mathcal{V}_i)}{|\mathcal{V}_i|}.$$

Note that there are also other graph partitioning objectives that fall under this framework (see, e.g., Section 3 of [3] which generalizes association and cut problems to weighted variants), and our approach can also be applied to these objectives.

2.2 Weighted kernel k-means and graph cuts

Consider the k-means objective function with clusters ℓ_1, \dots, ℓ_k :

$$\sum_{c=1}^k \sum_{\mathbf{x}_i \in \ell_c} \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2,$$

where $\boldsymbol{\mu}_c = (1/|\ell_c|) \cdot \sum_{\mathbf{x}_i \in \ell_c} \mathbf{x}_i$. It is straightforward to extend this to the weighted setting by introducing a weight w_i for each data point, which yields the following:

$$\sum_{c=1}^k \sum_{\mathbf{x}_i \in \ell_c} w_i \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2,$$

where now the mean $\boldsymbol{\mu}_c$ is the weighted mean $\boldsymbol{\mu}_c = \sum_{\mathbf{x}_i \in \ell_c} w_i \mathbf{x}_i / \sum_{\mathbf{x}_i \in \ell_c} w_i$. Further, we can replace the original data with mapped data $\phi(\mathbf{x}_i)$ and treat the entire problem in kernel space by expressing both the k-means algorithm, along with the objective, in terms of inner products. This is necessary for the connection to graph cuts.

Dhillon et al. [3] showed that there is a connection between the weighted kernel k-means objective and several spectral graph cut objectives. We will discuss in particular the connection to normalized cuts. Define the degree matrix D as the diagonal matrix whose entries D_{ii} are equal to the degree of node i . The surprising fact established in [3] is that normalized cuts and weighted kernel k-means are mathematically equivalent, in the following sense: if A is an adjacency matrix, then the normalized cuts objective on A is equivalent to the weighted kernel k-means objective (plus a constant) on the kernel matrix $K = \rho D^{-1} + D^{-1} A D^{-1}$, where ρ is chosen such that K is a positive semi-definite matrix, and where the weights of the data points are equal to the degrees of the nodes. Thus, for the purposes of minimizing the weighted (kernel) k-means objective function, we can effectively interchange the objective with the normalized cut objective, i.e.,

$$\min \sum_{c=1}^k \sum_{\mathbf{x} \in \ell_c} w_i \|\mathbf{x} - \boldsymbol{\mu}_c\|^2 \equiv \min \sum_{c=1}^k \frac{\text{cut}(\mathcal{V}_c, \mathcal{V}/\mathcal{V}_c)}{\text{deg}(\mathcal{V}_c)} \tag{1}$$

for the appropriate definition of the kernel matrix. In particular, this result gives an algorithm for monotonically minimizing the normalized cut objective—we just form the appropriate kernel and set the weights to the degrees, and then run weighted kernel k-means on that kernel matrix. Similar equivalences hold for both the ratio cut and ratio association objectives—by forming appropriate kernels and weights, the graph cut objectives can be shown to be mathematically equivalent to the weighted kernel k-means objective.

3 The Power-law Normalized Cut Objective

Our goal is to propose and study a new set of graph cut objectives that produce power-law distributed cluster sizes. In order to achieve this, we will borrow some key ideas from Bayesian nonparametrics. More specifically, we look at the Pitman-Yor Chinese restaurant process [8], a generalization of the Chinese Restaurant Process that specifically yields power-law distributed cluster sizes. For simplicity, we will focus on the normalized cut objective as an example. One can simply replace the normalized cut objective with other graph cut objectives to obtain other power-law graph cut algorithms in our framework.

3.1 Pitman-Yor EPPF

The canonical Bayesian nonparametric clustering prior is the Chinese restaurant process (CRP) [7]. It yields a distribution on clusterings such that the number of clusters are not fixed, and where the sizes of the clusters decay exponentially. The description of the CRP is as follows: customers enter a restaurant with an infinite number of tables (each table corresponds to a cluster). The first customer sits at the first table. Subsequent customers sit at tables with probability proportional to the number of seated customers at that table, and with probability proportional to α sit at a new table. The Pitman-Yor process leads to an extension of the CRP such that the cluster sizes instead follow a power-law distribution. In this modified version of the CRP, when customers sit down at tables, they sit at an existing table with probability proportional to the number of existing occupants minus θ ($0 \leq \theta < 1$), and at a new table with probability proportional to $k \cdot \theta + \alpha$, where k is the current number of occupied tables. Thus, as the number of tables k increases, there is a higher probability of starting a new table; this leads to the heavier power-law distribution of cluster sizes.

One can explicitly write down the probability of observing a particular seating arrangement under the Pitman-Yor CRP, and the resulting formula is known as the Pitman-Yor exchangeable partition probability

function (EPPF) [9]. If we let Z be an indicator matrix for the resulting clustering, then the probability distribution of Z under the Pitman-Yor CRP is expressed by the following unintuitive and somewhat cumbersome form:

$$p(Z|\alpha, \theta) = \frac{[\alpha + \theta]_{k-1, \theta}}{[\alpha + 1]_{n-1}} \cdot \prod_{c=1}^k [1 - \theta]_{n_c-1} \quad (2)$$

where

$$[x]_{m,a} = \begin{cases} 1 & m = 0 \\ x \cdot (x + a) \cdots (x + (m - 1)a) & m = 1, 2, \dots, \end{cases}$$

n_c is the size of cluster c , and $[x]_m$ is defined as $[x]_{m,1}$. One can verify that, when $\theta = 0$, we actually obtain the original CRP probability distribution. One can also show that the expected number of clusters under this distribution is $\mathcal{O}(\alpha n^\theta)$, and that we obtain the desired power-law cluster size distribution [9].

3.2 Power-law normalized cut objective

To obtain power-law distributed cluster sizes within a graph clustering setting, we treat the Pitman-Yor EPPF as a regularizer for the cluster indicator matrix of normalized cuts. Then our resulting objective is given as below:

$$\min_{Z,k} \sum_{c=1}^k \frac{\text{cut}(\mathcal{V}_c, \mathcal{V} \setminus \mathcal{V}_c)}{\text{deg}(\mathcal{V}_c)} + \lambda \cdot r_{\alpha, \theta}(Z), \quad (3)$$

$$r_{\alpha, \theta}(Z) = -\ln p(Z|\alpha, \theta)$$

where $Z = [z_1, \dots, z_n]$ is the indicator for the cluster assignment of each node, $r_{\alpha, \theta}(Z)$ is the negative log of the Pitman-Yor EPPF and λ is a tradeoff between the original graph cut objective and the regularization term. The first term is the standard normalized cut objective. The desired power-law distributed partition would give a high value of the Pitman-Yor EPPF and thus a low value of the second term. Therefore, the clustering result that minimizes this objective should give a partition of the graph such that both similarity information is preserved and cluster sizes are power-law distributed.

4 Optimization

The objective function (3) defined in the previous section enforces a tradeoff between standard normalized cuts and a preference for power-law cluster size structure. We now turn to optimization of the resulting objective.

4.1 The vector case

The first observation that we can make is that spectral methods will not apply to our proposed objective. Recall that for the normalized cut objective, a standard

approach is to relax the cluster indicator matrix Z to be continuous, leading to a simple eigenvector problem that can be optimized globally. When applying such a technique to the power-law normalized cut objective, one would need to incorporate the regularization term appropriately into the trace maximization problem that emerges from the spectral solution, but this turns out to be impossible.

Instead we must turn to the other main optimization strategy for normalized cuts—namely the equivalence to weighted kernel k-means—and we will adapt the weighted kernel k-means algorithm for our problem. To start, in this section we will derive a k-means-like algorithm for the following regularized k-means problem:

$$\min_{Z,k} \sum_{c=1}^k \sum_{\mathbf{x}_i \in \ell_c} w_i \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2 + \lambda \cdot r_{\alpha, \theta}(Z),$$

where the means $\boldsymbol{\mu}_c$ are the weighted means of the points in ℓ_c as in standard weighted k-means as discussed in Section 2.2. Once we have obtained the algorithm for this case, we can easily extend the connection between normalized cuts and weighted kernel k-means to obtain an algorithm for monotonic local convergence of the power-law normalized cut objective. Note that this treatment is equally applicable to the ratio cut and ratio association objectives.

We observe that, when the cluster indicators Z are fixed, the weighted mean is justified in the above objective since it is the best cluster representative for each cluster in terms of the objective function, i.e., for fixed Z and any choice of c , the regularizer is constant and we have by simple differentiation

$$\sum_{\mathbf{x}_i \in \ell_c} w_i \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2 = \min_{\mathbf{m}} \sum_{\mathbf{x}_i \in \ell_c} w_i \|\mathbf{x}_i - \mathbf{m}\|^2.$$

Therefore, the updates on $\boldsymbol{\mu}_c$ will be exactly as in standard weighted k-means.

The other step is the update on the indicators Z . In standard k-means, these updates are derived by fixing the means and minimizing the k-means objective function with respect to each z_i , which yields the usual k-means assignment step. The Pitman-Yor EPPF regularizer makes the assignment updates somewhat less trivial, but it is still fairly straightforward. For each data point we consider the objective function when assigning that point to every existing cluster, as well as to a new cluster, and assign to the cluster that results in the smallest objective function. The regularizer effectively adds a “correction” to each distance computation $w_i \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2$. Let n_c be the number of points in cluster c . After going through the algebra, we arrive at the following: if \mathbf{x}_i is currently assigned to ℓ_c , then we

have that the distance to another cluster $\ell_{c'}$ (ignoring constants, which do not affect re-assignment) is given by the following cases:

- If $c' = c$:

$$d(\mathbf{x}_i, c') = w_i \|\mathbf{x}_i - \boldsymbol{\mu}_{c'}\|^2.$$

- if $n_c > 1$ and c' is an existing cluster,

$$d(\mathbf{x}_i, c') = w_i \|\mathbf{x}_i - \boldsymbol{\mu}_{c'}\|^2 + \lambda \cdot \ln \left(\frac{n_c - 1 - \theta}{n_{c'} - \theta} \right).$$

- if $n_c = 1$ and c' is an existing cluster,

$$d(\mathbf{x}_i, c') = w_i \|\mathbf{x}_i - \boldsymbol{\mu}_{c'}\|^2 + \lambda \cdot \ln \left(\frac{\alpha + (k-1)\theta}{n_{c'} - \theta} \right).$$

- if $n_c > 1$ and c' is a new cluster

$$d(\mathbf{x}_i, c') = \lambda \cdot \ln \left(\frac{n_c - 1 - \theta}{\alpha + k\theta} \right).$$

- if $n_c = 1$ and c' is a new cluster

$$d(\mathbf{x}_i, c') = \infty.$$

Observe that the distance to new clusters goes down as k increases, which is analogous to the property in the Pitman-Yor version of the Chinese restaurant process of being more likely to start a new table as the number of tables increases. In a similar way, when computing the distance to existing clusters, the distance becomes smaller as the cluster gets larger (i.e., as $n_{c'}$ goes up), leading to the “rich gets richer” behavior. Finally, whenever a new cluster is started by some point \mathbf{x}_i , we immediately set the mean to be \mathbf{x}_i . See Algorithm 1 for a full specification. Note that, analogous to the convergence proof of k-means, one can easily show that this algorithm monotonically decreases the regularized k-means objective until local convergence.

4.2 Power-law normalized cut algorithm

Recall that in Section 2.2 we discussed the equivalence between the graph cuts formulation and the weighted kernel k-means objective, as in (1). With this equivalence in hand, the extension from the vector case to the power-law graph cut objectives follows easily: we simply replace the weighted k-means term with a graph cuts term, which gives exactly the same objective with our power-law graph cuts objective in (3) up to a constant; then we apply Algorithm 1 in kernel space to solve the resulting optimization problem.

More specifically, given a graph $G = (\mathcal{V}, \mathcal{E})$ with adjacency matrix A , our power-law normalized cut algorithm is described as follows:

Algorithm 1 Power-law-means (vector case)

Input: $\mathbf{x}_1, \dots, \mathbf{x}_n$: data points; w_1, \dots, w_n : weights; λ : trade-off parameter; α, θ : Pitman-Yor EPPF parameters

Output: Clustering ℓ_1, \dots, ℓ_k ; k : number of clusters
 1: Init. $k = 1, \ell_1 = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}, \boldsymbol{\mu}_1$ the global mean
 2: Init. cluster indicators $z_i = 1$ for all $i = 1, \dots, n$.
 3: Repeat 4 to 17 until convergence.
 4: **for** each data point \mathbf{x}_i , suppose \mathbf{x}_i is currently assigned to cluster ℓ_c **do**

5: **if** $n_c = 1$, i.e. \mathbf{x}_i is a singleton cluster **then**
 6: compute its “regularized” distance $d(\mathbf{x}_i, c')$ to the other clusters according to the following:

- If $c' = c$, $d(\mathbf{x}_i, c') = 0$.
- if $c' \neq c$ and $\ell_{c'}$ is an existing cluster, $d(\mathbf{x}_i, c') = w_i \|\mathbf{x}_i - \boldsymbol{\mu}_{c'}\|^2 + \lambda \cdot \ln \left(\frac{\alpha + (k-1)\theta}{n_{c'} - \theta} \right)$
- if $\ell_{c'}$ is a new cluster, $d(\mathbf{x}_i, c') = \infty$

7: **else**

8: compute its “regularized” distance $d(\mathbf{x}_i, c')$ to the other clusters according to the following:

- If $c' = c$, $d(\mathbf{x}_i, c') = w_i \|\mathbf{x}_i - \boldsymbol{\mu}_{c'}\|^2$.
- if $c' \neq c$ and $\ell_{c'}$ is an existing cluster, $d(\mathbf{x}_i, c') = w_i \|\mathbf{x}_i - \boldsymbol{\mu}_{c'}\|^2 + \lambda \cdot \ln \left(\frac{n_c - 1 - \theta}{n_{c'} - \theta} \right)$
- if $\ell_{c'}$ is a new cluster, $d(\mathbf{x}_i, c') = \lambda \cdot \ln \left(\frac{n_c - 1 - \theta}{\alpha + k\theta} \right)$

9: **end if**

10: Assign \mathbf{x}_i to the cluster corresponding to the smallest regularized distance. Update Z :

$$z_i = \arg \min_{c'} d(\mathbf{x}_i, c').$$

11: **if** z_i corresponds to a new cluster **then**

12: set $k \leftarrow k + 1, z_i = k$, and $\boldsymbol{\mu}_k = \mathbf{x}_i$.

13: **end if**

14: **end for**

15: **for** each cluster ℓ_c **do**

16: Update $\boldsymbol{\mu}_c$ based on the weighted mean of the data points in cluster ℓ_c :

$$\boldsymbol{\mu}_c = \frac{\sum_{\mathbf{x} \in \ell_c} w_i \mathbf{x}_i}{\sum_{\mathbf{x} \in \ell_c} w_i}.$$

17: **end for**

- Compute the degree matrix D from A as the diagonal matrix whose entries D_{ii} are equal to the degree of node i .
- Compute the kernel matrix K from A using $K =$

$$\rho D^{-1} + D^{-1} A D^{-1}.$$

- Run Algorithm 1 in kernel space with kernel K and weights given by the degrees of the nodes. The power-law normalized cut clustering result is then obtained directly from Algorithm 1.

In kernel space, the regularized distance remains unchanged. The only change is that now we need to compute $\|\phi(\mathbf{x}_i) - \boldsymbol{\mu}_c\|^2$ instead of $\|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2$. We expand the last distance computation and use the formula for $\boldsymbol{\mu}_c$ and obtain:

$$\begin{aligned} \phi(\mathbf{x}_i) \cdot \phi(\mathbf{x}_i) &- \frac{2 \sum_{\mathbf{x}_j \in \ell_c} w_j \phi(\mathbf{x}_i) \cdot \phi(\mathbf{x}_j)}{\sum_{\mathbf{x}_j \in \ell_c} w_j} \\ &+ \frac{\sum_{\mathbf{x}_j, \mathbf{x}_l \in \ell_c} w_j w_l \phi(\mathbf{x}_j) \cdot \phi(\mathbf{x}_l)}{(\sum_{\mathbf{x}_j \in \ell_c} w_j)^2}. \end{aligned}$$

Using the kernel matrix K , the above may be written as:

$$K_{ii} - \frac{2 \sum_{\mathbf{x}_j \in \ell_c} w_j K_{ij}}{\sum_{\mathbf{x}_j \in \ell_c} w_j} + \frac{\sum_{\mathbf{x}_j, \mathbf{x}_l \in \ell_c} w_j w_l K_{jl}}{(\sum_{\mathbf{x}_j \in \ell_c} w_j)^2}.$$

We note that, as when applying weighted kernel k-means to the standard normalized cut problem [3], each iteration of Algorithm 1 when applied in kernel space with K requires time $O(|\mathcal{E}|)$, making it very scalable for applications to large graphs. Also note that by using an appropriate kernel matrix K , we can utilize other graph cut objectives in this framework.

4.3 Connection to Pitman-Yor MAP inference

Finally, we briefly consider the connections between our proposed objective and a simple Pitman-Yor process mixture model. Consider the following Bayesian nonparametric generative model:

$$\begin{aligned} Z &\sim \text{PYCRP}(\alpha, \theta); \\ \mathbf{x}_i &\sim \mathcal{N}(\boldsymbol{\mu}_{z_i}, (\sigma/(2w_i))I), \quad i = 1, \dots, n, \end{aligned}$$

where PYCRP refers to the Pitman-Yor Chinese Restaurant Process. To perform MAP inference, we can write down the joint likelihood and maximize with respect to the relevant parameters:

$$\begin{aligned} &\text{argmax}_{Z, k, \boldsymbol{\mu}} p(X, Z) \\ &\equiv \text{argmin}_{Z, k, \boldsymbol{\mu}} - \ln \left(\prod_{i=1}^n \mathcal{N}(\boldsymbol{\mu}_{z_i}, \frac{\sigma}{2w_i} I) \cdot p(Z | \alpha, \theta) \right) \\ &\equiv \text{argmin}_{Z, k, \boldsymbol{\mu}} \frac{1}{\sigma} \sum_{c=1}^k \sum_{\mathbf{x}_i \in \ell_c} w_i \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2 - \ln p(Z | \alpha, \theta) \\ &\equiv \text{argmin}_{Z, k, \boldsymbol{\mu}} \sum_{c=1}^k \sum_{\mathbf{x} \in \ell_c} w_i \|\mathbf{x}_i - \boldsymbol{\mu}_c\|^2 + \lambda \cdot r_{\alpha, \theta}(Z), \end{aligned}$$

where $\lambda = \sigma$. Note that the minimization with respect to $\boldsymbol{\mu}$ yields precisely the weighted means, and so based on the equivalence between weighted kernel k-means and normalized cuts, we can see that our proposed objective function may be viewed in a MAP inference framework. This framework also justifies the use of the log of the Pitman-Yor EPPF as a regularizer.

4.4 Comparison to existing power-law clustering algorithm *pyp-means*

In [12], the authors propose a different objective for power-law data clustering, namely:

$$\text{arg min}_{l_1, \dots, l_n} \sum_{c=1}^k \sum_{i \in \ell_c} \|\mathbf{x}_i - \boldsymbol{\mu}_c\| + (\lambda - \ln k \cdot \theta)k,$$

which adds a $-k \ln k \cdot \theta$ term to the *dp-means* objective function [10]. While this objective function does incorporate the number of clusters into the optimization, it does not require or encourage the cluster sizes to follow a power-law distribution. Moreover, in their experiments, the authors set $\theta = \lambda/6$. In this case, the objective function becomes:

$$\text{arg min}_{l_1, \dots, l_n} \sum_{c=1}^k \sum_{i \in \ell_c} \|\mathbf{x}_i - \boldsymbol{\mu}_c\| + \lambda \left(1 - \frac{\ln k}{6} \right) k,$$

One can show that, when the number of data points exceeds $\ln 403 \approx 6$, the trivial clustering result, namely every data point is a singleton cluster, will minimize this objective. This can be seen by the fact that the trivial clustering result minimizes the k -means objective by simply being 0 and that $k =$ the number of data points minimizes the regularization term. In short, this objective is not appropriate for power-law clustering applications. In the following experiment section, we will also compare our algorithm with their method empirically.

5 Experiments

We conclude with a brief set of experiments demonstrating the utility of our methods. Namely, we will show that our approach enjoys benefits over the k-means algorithm on real power-law datasets in the vector setting and benefits over standard normalized cuts¹ on synthetic and real data in the graph setting. We also compare our method with the *pyp-means* [12] and show that our method achieves better clustering results. Throughout the experiments, we use normalized mutual information (NMI) between the algorithm's clusters and the ground-truth clusters for evaluation.

¹Normalized cut image segmentation code: <http://www.cis.upenn.edu/~jshi/software/>.




Figure 1: Results on a Pitman-Yor generated stochastic block model graph. Left: Adjacency matrix of the graph, indexed by clusters. Middle: Power-law normalized cut results; NMI is 0.866. Right: Normalized Cuts result; NMI is 0.687.




Figure 2: Cluster size distributions on the *ecoli* data set. Left: Ground-truth. Middle left: Algorithm 1; NMI is 0.723. Middle Right: *pyp*-means; NMI is 0.608. Right: Result of *k*-means; NMI is 0.582.

Synthetic power-law graph data. We begin with a synthetic power-law random graph dataset generated by the Pitman-Yor process applied to the stochastic block model. Specifically, the Pitman-Yor CRP is first used to generate data cluster assignments and then a standard stochastic block model uses the assignments to generate a random graph. We create a dataset with 4000 nodes with 14 disjoint clusters using the above process, with the corresponding adjacency matrix shown in the left of Figure 1. The parameters α and θ we use in the Pitman-Yor process model are 1 and 0.2, respectively. In the stochastic block model, the stochastic block matrix is sampled from two Gaussian distributions: one being $\mathcal{N}(0.3, 0.001)$ for diagonal entries and the other being $\mathcal{N}(0.01, 0.001)$ for non-diagonal entries. Our power-law normalized cut algorithm is then applied on this dataset with parameters validated on a separate validation dataset generated from the same process. We compare with normalized cuts with its k set to be the ground-truth. The results are shown in Figure 1; normalized cuts splits the big clusters while our algorithm nearly produces the ground-truth clusters.

Real world power-law data sets. Next we consider comparing Algorithm 1 with *k*-means and *pyp*-means [12] on real world benchmark data sets to demonstrate that our algorithm performs best on clustering vector data when cluster sizes are power-law distributed. We selected 6 UCI classification datasets whose class labels are power-law distributed (see Fig-

Table 1: NMI scores on a set of UCI power-law datasets.

Dataset	NMI		
	Ours	<i>k</i> -means	<i>pyp</i> -means
<i>audiology</i>	0.621	0.518	0.417
<i>ecoli</i>	0.700	0.545	0.608
<i>glass</i>	0.427	0.315	0.297
<i>hypothyroid</i>	0.024	0.009	0.077
<i>page-blocks</i>	0.209	0.123	0.088
<i>flags</i>	0.275	0.198	0.178

ure 2) and use class labels as the ground-truth for clusters². Each dataset is then randomly split 30/70 for validation/clustering. We normalize the datasets so that the values of all features lie in $[0, 1]$. On each validation set, we validate the parameters of Algorithm 1 (i.e. λ, α, θ) and the parameters of *pyp*-means only to yield cluster numbers close to the ground-truth (to make a fair comparison with *k*-means). On each clustering set, we use the validated parameter settings for Algorithm 1 and *pyp*-means and the ground-truth k for *k*-means to perform the clustering. The NMI are computed between the ground-truth and the computed clusters, and results are averaged over 10 runs, as shown in Table 1. As we can see, Algorithm 1 per-

²See <https://archive.ics.uci.edu/ml/datasets.html> for detailed descriptions of the selected data sets.

Table 2: NMI scores on graphs generated from UCI power-law datasets.

Dataset	NMI	
	Ours	Normalized cuts
<i>audiology</i>	0.662	0.561
<i>ecoli</i>	0.702	0.591
<i>glass</i>	0.432	0.356
<i>hypothyroid</i>	0.011	0.008
<i>page-blocks</i>	0.222	0.126
<i>flags</i>	0.357	0.200

forms better than k-means on all 6 datasets in terms of NMI. Also, it is better than *pyp*-means on all datasets except on the *hypothyroid*. Note that the *pyp*-means is better than *k*-means in 3 datasets and worse than *k*-means in the other 3. Such high variance results on power-law datasets make us doubt that *pyp*-means is really able to achieve power-law clustering. In Figure 2, we show the resulting clusterings on the *ecoli* dataset given by Algorithm 1, *pyp*-means and k-means, in which we use the whole dataset for clustering with validated parameters. It is clear that k-means produces more uniform clusters and *pyp*-means also splits the largest cluster in the dataset.

Real world power-law graph data sets. In this part we convert the UCI vector datasets used in the preceding experiment to form power-law graphs and perform power-law normalized cuts on these graphs. We also run normalized cuts algorithm on these graphs to compare with our method. To obtain the graphs, we first form the adjacency matrix by using a Gaussian similarity kernel on the vector data after normalizing them to $[0, 1]$. Then we use the adjacency matrix to form the kernel matrix and the weights as discussed in Section 4.2. We randomly split data into validation/clustering with ratio of 30/70. Parameters are selected on the validation set so that cluster numbers are close to the ground-truth. The number of clusters in normalized cuts is set to the true number of clusters. This is again for a fair comparison with normalized cuts. Finally, we apply our power-law normalized cuts and normalized cuts on the clustering dataset. NMI averaged over 10 runs are shown in Table 2. As we can see, our power-law normalized cuts is better than normalized cuts on all the graphs in terms of NMI.

Image segmentation. Finally, we briefly demonstrate some qualitative results on image segmentation on the Berkeley segmentation data set [16]. We adopt an approach that is similar to the approach in [1] to compute the affinity matrix. Then we perform our power-law normalized cuts with the affinity matrix.




Figure 3: Image segmentation results. Left to right: ground-truth, our method, normalized cuts. This figure is best viewed in color.

We compare standard normalized cuts with our proposed method on graphs generated from input images. Figure 3 displays some example images; we see that normalized cuts tends to break up large segments more often than our approach.

6 Conclusion

We proposed a general framework of power-law graph cut algorithms that produce clusters whose sizes are power-law distributed, and also does not fix the number of clusters upfront. The Pitman-Yor exchangeable partition probability function (EPPF) was incorporated into power-law graph cut objectives as a regularizer to promote power-law cluster size distributions. A simple iterative algorithm was then proposed to locally optimize several objectives. Our proposed algorithm can be viewed as performing MAP inference on a particular Pitman-Yor mixture model. Finally, we conducted experiments on various data sets and compared to existing baselines.

Acknowledgements

This work was supported by NSF award IIS-1217433.

References

- [1] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8):888–905, 2000.
- [2] Pak K Chan, Martine DF Schlag, and Jason Y Zien. Spectral k-way ratio-cut partitioning and clustering. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 13(9):1088–1096, 1994.
- [3] Inderjit S Dhillon, Yuqiang Guan, and Brian Kulis. Weighted graph cuts without eigenvectors a multilevel approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):1944–1957, 2007.
- [4] DM Greig, BT Porteous, and Allan H Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279, 1989.
- [5] Erik B Sudderth and Michael I Jordan. Shared segmentation of natural scenes using dependent Pitman-Yor processes. In *NIPS*, pages 1585–1592, 2008.
- [6] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
- [7] Nils Hjort, Chris Holmes, Peter Mueller, and Stephen Walker. *Bayesian Nonparametrics: Principles and Practice*. Cambridge University Press, Cambridge, UK, 2010.
- [8] Jim Pitman and Marc Yor. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, pages 855–900, 1997.
- [9] Jim Pitman. *Combinatorial Stochastic Processes*. Springer-Verlag, 2006. Lectures from the Saint-Flour Summer School on Probability Theory.
- [10] Brian Kulis and Michael I Jordan. Revisiting k-means: New algorithms via Bayesian nonparametrics. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 513–520, 2012.
- [11] Tamara Broderick, Brian Kulis, and Michael I Jordan. MAD-bayes: MAP-based asymptotic derivations from Bayes. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, 2013.
- [12] Xuhui Fan, Yiling Zeng, and Longbing Cao. Non-parametric power-law data clustering. *CoRR*, abs/1306.3003, 2013.
- [13] Kevin Lang. Fixing two weaknesses of the spectral method. In *Advances in Neural Information Processing Systems*, pages 715–722, 2005.
- [14] Brian W Kernighan and Shen Lin. An efficient heuristic procedure for partitioning graphs. *Bell system technical journal*, 49(2):291–307, 1970.
- [15] Stella X Yu and Jianbo Shi. Multiclass spectral clustering. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 313–319. IEEE, 2003.
- [16] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(5):898–916, 2011.