Guarantees for Spectral Clustering with Fairness Constraints

Matthäus Kleindessner, Samira Samadi, Pranjal Awasthi, Jamie Morgenstern

RUTGERS | Georgia Tech

Spectral clustering (SC)

SC is the method of choice for clustering the nodes of a graph.

Example: want to find clusters in a friendship network — many connections within clusters, few connections between clusters

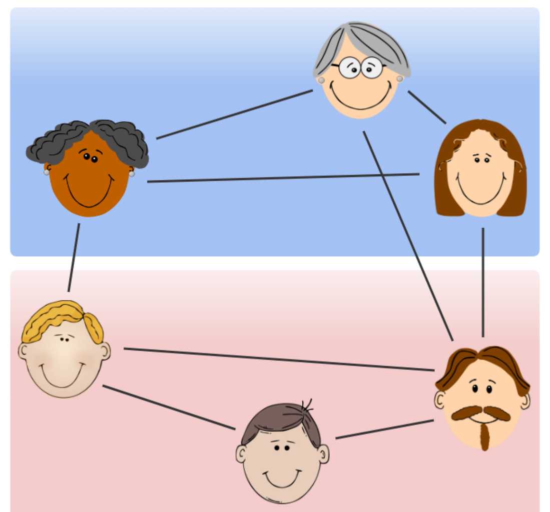

Friendship network and a clustering into two clusters obtained from running SC. According to the fairness notion of Chierichetti et al. (2017) this clustering is highly unfair with respect to the two demographic groups of males and females (see below).

There are several versions of SC (von Luxburg, 2007); here we focus on *unnormalized SC*, but our results hold similarly for *normalized SC* too.

Formally:

$G=(V,E)$ ... undirected graph on $V=[n]$; $W\in\mathbb{R}^{n\times n}$ ... (weighted) adjacency matrix with $W_{ij}>0$ iff there is an edge between i and j; k ... number of clusters

SC aims to partition V into k clusters with minimum value of the *RatioCut objective function* defined as follows: for a clustering $V=C_1\dot\cup\ldots\dot\cup C_k$ it is

$$\mathrm{RatioCut}(C_1,\ldots,C_k)=\sum_{l=1}^{k}\frac{\mathrm{Cut}(C_l,V\setminus C_l)}{|C_l|}\quad\text{with}\quad \mathrm{Cut}(C_l,V\setminus C_l)=\sum_{\substack{i\in C_l,\\ j\in V\setminus C_l}}W_{ij}.$$

Let D be the degree matrix (with $D_{ii}=\sum_{j\in[n]}W_{ij}$) and $L=D-W$ be the *graph Laplacian matrix*. Encoding a clustering $V=C_1\dot\cup\ldots\dot\cup C_k$ by $H\in\mathbb{R}^{n\times k}$ with

$$H_{il}=\begin{cases}1/\sqrt{|C_l|}, & i\in C_l,\\ 0, & i\notin C_l\end{cases},\tag{1}$$

we have $\mathrm{RatioCut}(C_1,\ldots,C_k)=\mathrm{Tr}(H^TLH)$. SC relaxes the exact problem

$$\min_{H\in\mathbb{R}^{n\times k}}\mathrm{Tr}(H^TLH)\quad\text{subject to } H \text{ is of form (1)}\tag{2}$$

and solves

$$\min_{H\in\mathbb{R}^{n\times k}}\mathrm{Tr}(H^TLH)\quad\text{subject to } H^TH=I_k\tag{3}$$

instead. A solution to (3) is given by a matrix H that contains some orthonormal eigenvectors corresponding to the k smallest eigenvalues of L as columns. Then, a clustering of V is inferred from H by applying k-means clustering to the rows of H.

Spectral clustering with fairness constraints

Let V contain h demographic groups V_s such that $V=\dot\cup_{s\in[h]}V_s$. Chierichetti et al. (2017) proposed a notion of fairness for clustering: **in every cluster, each group V_s should be represented with (approximately) the same fraction as in the whole data set V.**

We can incorporate this goal into (2) via a linear constraint on H:

Algorithm 1 Unnormalized SC with fairness constraints

Input: adj. matrix $W\in\mathbb{R}^{n\times n}$; $k\in\mathbb{N}$; group-membership vectors $f^{(s)}\in\{0,1\}^n$

Output: a clustering of $[n]$ into k clusters

- compute the Laplacian matrix $L=D-W$
- let F be a matrix with columns $f^{(s)}-\frac{|V_s|}{n}\cdot\mathbf{1}_n$, $s\in[h-1]$
- let the columns of Z form an orthonormal basis of the nullspace of F^T
- compute the k smallest (respecting multiplicities) eigenvalues of Z^TLZ and the corresponding orthonormal eigenvectors (written as columns of Y)
- apply k-means clustering to the rows of $H=ZY$

Lemma 1 *For $s\in[h]$, let $f^{(s)}\in\{0,1\}^n$ be the group-membership vector of V_s, that is $f_i^{(s)}=1$ if $i\in V_s$ and $f_i^{(s)}=0$ otherwise. Let $V=C_1\dot\cup\ldots\dot\cup C_k$ be a clustering that is encoded as in (1). We have, for every $l\in[k]$,*

$$\forall s\in[h-1]:\sum_{i=1}^{n}\left(f_i^{(s)}-\frac{|V_s|}{n}\right)H_{il}=0\quad\Leftrightarrow\quad\forall s\in[h]:\frac{|V_s\cap C_l|}{|C_l|}=\frac{|V_s|}{n}.$$

Let $F\in\mathbb{R}^{n\times(h-1)}$ comprise the vectors $f^{(s)}-(|V_s|/n)\cdot\mathbf{1}_n$, $s\in[h-1]$, as columns. According to Lemma 1, instead of (3) we should solve

$$\min_{H\in\mathbb{R}^{n\times k}}\mathrm{Tr}(H^TLH)\quad\text{subject to } H^TH=I_k \text{ and } F^TH=0_{(h-1)\times k}.\tag{4}$$

Our approach is analogous to existing versions of constrained SC that try to incorporate must-link constraints into the SC framework (e.g., Yu and Shi, 2004).

How to solve (4)? Let $Z\in\mathbb{R}^{n\times(n-h+1)}$ be a matrix whose columns form an orthonormal basis of the nullspace of F^T. We can substitute $H=ZY$ for $Y\in\mathbb{R}^{(n-h+1)\times k}$, and then we end up with a problem of the form (3).

Our strategy is summarized in Algorithm 1. Note that, in general, there is no guarantee on how fair the output of Algorithm 1 is. However, often it is quite fair:

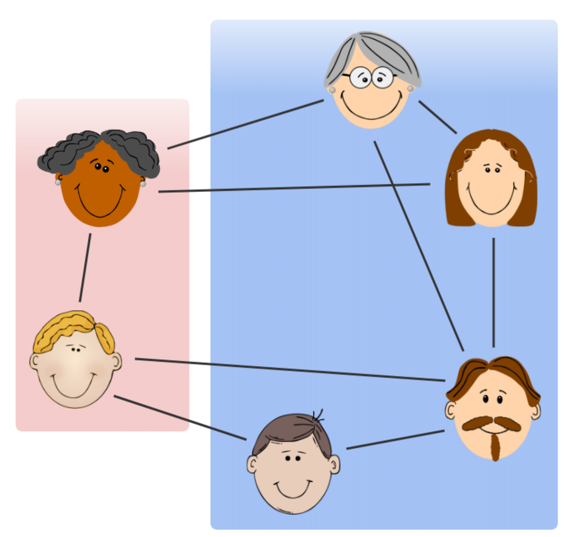

A clustering into two clusters obtained from running Algorithm 1. According to the fairness notion of Chierichetti et al. (2017) this clustering is perfectly fair.

Analysis on variant of the stochastic block model (SBM)

We analyze our fair version of SC on a natural variant of the famous SBM (Holland et al., 1983): let $V=[n]$ comprise h groups $V=\dot\cup_{s\in[h]}V_s$ and be partitioned into k ground-truth clusters $V=C_1\dot\cup\ldots\dot\cup C_k$ such that $|V_s\cap C_l|/|C_l|=\eta_s$, $s\in[h], l\in[k]$, for some $\eta_s\in(0,1)$ ($\Rightarrow$ the ground-truth clustering is perfectly fair). We define a random graph on V by connecting vertices i and j with probability $\Pr(i,j)$, where

$$\Pr(i,j)=\begin{cases}a, & i \text{ and } j \text{ in same cluster and in same group,}\\ b, & i \text{ and } j \textit{ not} \text{ in same cluster, but in same group,}\\ c, & i \text{ and } j \text{ in same cluster, but } \textit{not} \text{ in same group,}\\ d, & i \text{ and } j \textit{ not} \text{ in same cluster and } \textit{not} \text{ in same group}\end{cases}\tag{5}$$

for some $a>b>c>d$. An example of such a graph can be seen in the next figure.

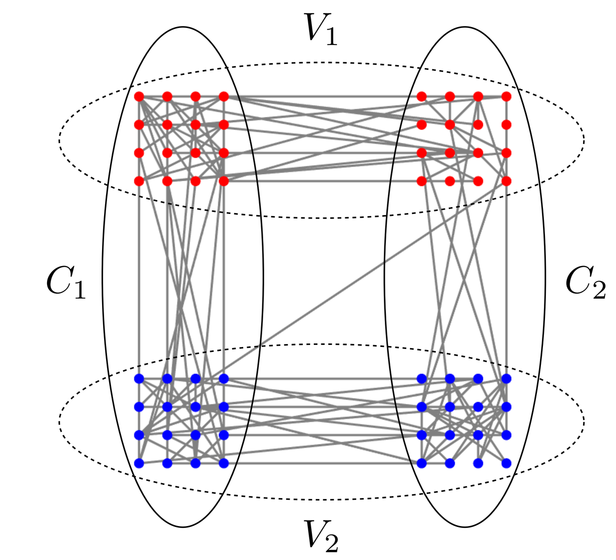

Example of a graph generated from our variant of the SBM. There are two meaningful clusterings into two clusters: $V=C_1\dot\cup C_2$ and $V=V_1\dot\cup V_2$. However, only the first one is fair.

For such a graph, standard SC is likely to return the unfair clustering $V=V_1\dot\cup V_2$ while Algorithm 1 returns the fair clustering $V=C_1\dot\cup C_2$ with high probability:

Theorem 1 *Let $V=[n]$ comprise h groups $V=\dot\cup_{s\in[h]}V_s$ and be partitioned into k ground-truth clusters $V=C_1\dot\cup\ldots\dot\cup C_k$ such that*

$$|V_s|=\frac{n}{h},\quad |C_l|=\frac{n}{k},\quad \frac{|V_s\cap C_l|}{|C_l|}=\frac{1}{h},\quad s\in[h],l\in[k].\tag{6}$$

Let G be a random graph constructed according to our variant of the SBM (5) with $a>b>c>d$ and $a\geq C\ln n/n$ for some $C>0$. Assume that we run Algorithm 1 on G, where we apply a $(1+M)$-approximation algorithm to the k-means problem encountered in the last step, for some $M>0$. Then there exist $\widehat{C},\widetilde{C}>0$ such that the following is true: if $a(c-d)^{-2}k^3n^{-1}\ln n<\widehat{C}(1+M)^{-1}$, then with probability at least $1-n^{-1}$, the clustering returned by Algorithm 1 "misclassifies" at most $\widetilde{C}a(c-d)^{-2}k^2(1+M)\ln n$ many vertices.

Experiments

SBM

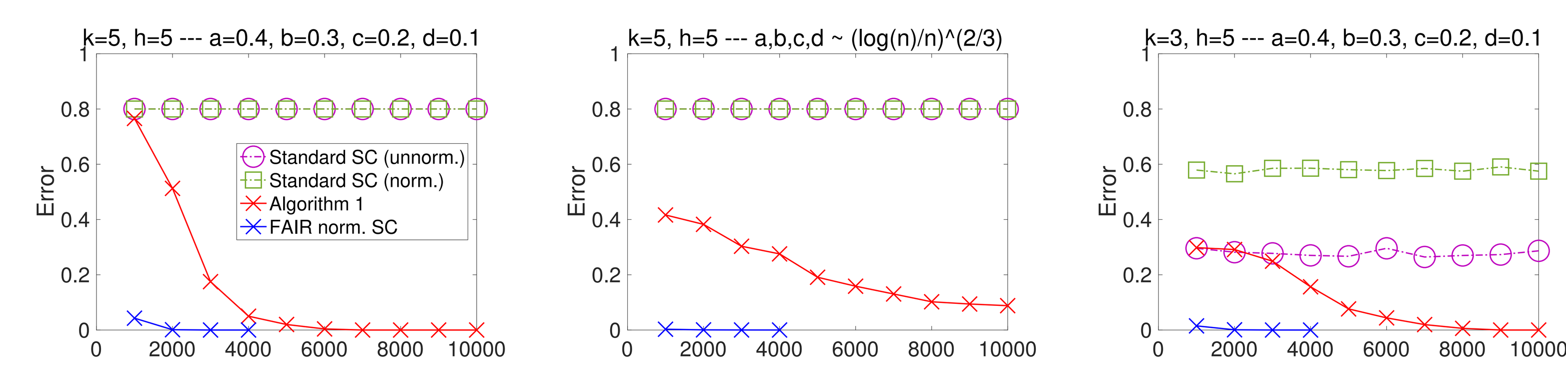

Error (fraction of "misclassified" vertices) as a function of n. In the third plot, C_1, C_2 and C_3 have different sizes and Assumption (6) is not satisfied.

Real networks

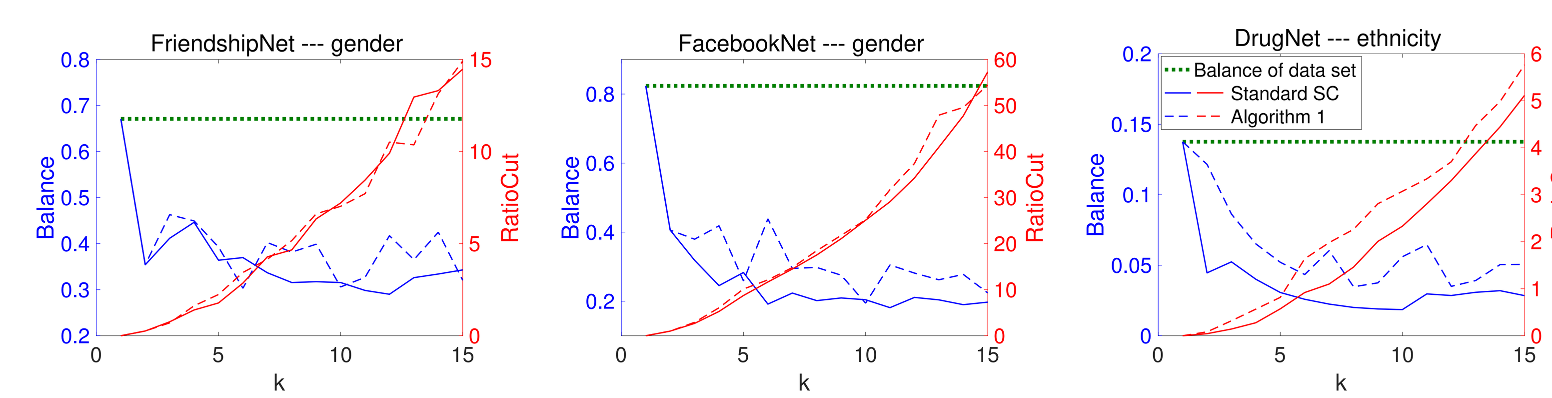

Average balance of clusters and RatioCut value as a function of k.

References

F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *Neural Information Processing Systems (NIPS)*, 2017.

P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5:109–137, 1983.

U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

S. X. Yu and J. Shi. Segmentation given partial grouping constraints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):173–183, 2004.

The face images were found on `https://openclipart.org` and are in the public domain.