Lecture VIII: Spectral Clustering

Marina Meilă
mmp@stat.washington.edu

Department of Statistics
University of Washington

March 2022

1. Graph clustering paradigms
2. MNCut and Random Walks
3. Embedding, the spectral mapping, lumpability
4. The algorithm
   - Affinity propagation

Similarity based clustering

- **Paradigm:** the features we observe are measures of **similarity/dissimilarity** between pairs of data points, e.g

	points	features
Image segmentation	pixels	distance in color space or location, separated by a contour, belong to same texture
Social network	people	friends, coworkers, phone calls, emails
Text analysis	words	appear in same context

- The features are summarized by a single **similarity measure** S_{ij}
  - e.g $S_{ij} = e^{\sum_k \alpha_k \text{feature}_k(i,j)}$ for all points i, j
  - symmetric $S_{ij} = S_{ji}$
  - non-negative $S_{ij} \geq 0$
- We want to put points that are similar to each other in the same cluster, dissimilar points in different clusters
- Problem is often cast as a **graph cut** problem
  - points = graph nodes, similarity S_{ij} = weight of edge ij
  -

Paradigms for grouping

- Graph cuts
  remove some edges $\Longrightarrow$ disconnected graph
  the groups are the connected components
- By "similar behavior"
  nodes i, j in the same group iff i, j "have the same pattern of connections" w.r.t other nodes
- By Embedding
- map nodes $V = \{1, 2, \ldots, n\} \longrightarrow \{x_1, x_2, \ldots, x_n\} \in \mathbb{R}^d$ then use standard classification and clustering methods

Definitions

- $V = \{1, 2, \ldots, n\}$
- node **degree** or **volume**

$$D_i = \sum_{j \in V} S_{ij}$$

- **volume** of cluster $C \subseteq V$

$$D_C = \sum_{i \in C} D_i$$

- **cut** between subsets $C, C' \subseteq V$

$$\sum_{i \in C} \sum_{j \in C'} S_{ij}$$

- **Multiway Normalized Cut** of a partition $\Delta = \{C_{1:K}\}$ of V

$$MNCut(\Delta) = \sum_{k=1}^{K} \sum_{k' \neq k} \frac{Cut(C_k, C_{k'})}{D_{C_k}}$$

in particular, for $K = 2$,

$$MNCut(C, C') = Cut(C, C') \left(\frac{1}{D_C} + \frac{1}{D_{C'}} \right)$$

Motivation for MNCut

$S_{ij} \propto 1/\text{dist}(i,j)$

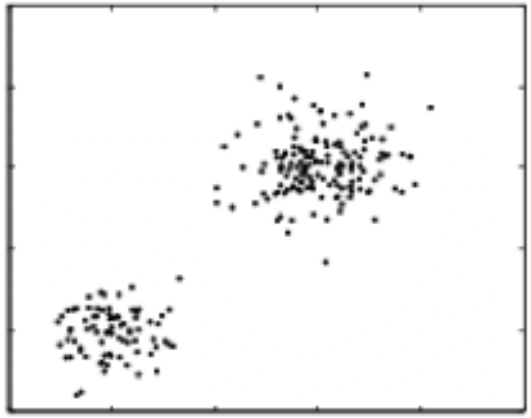

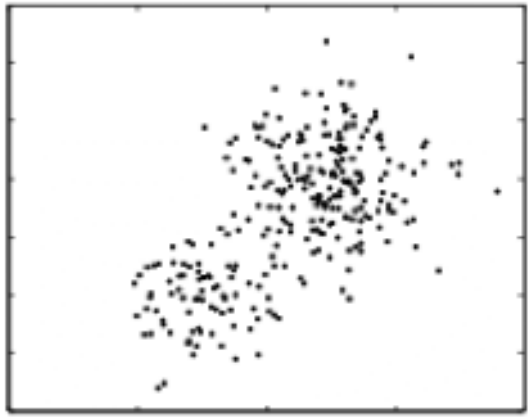

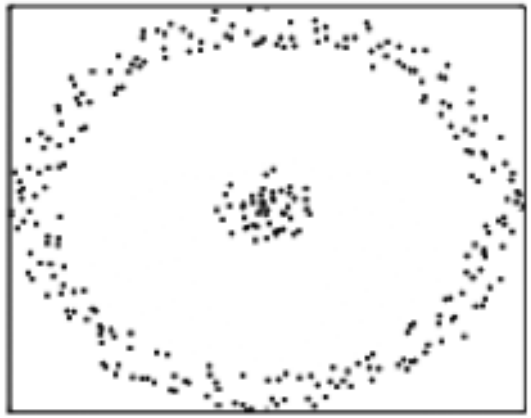

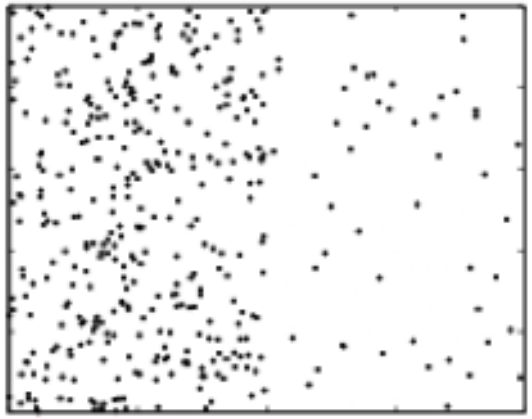

A random walks view

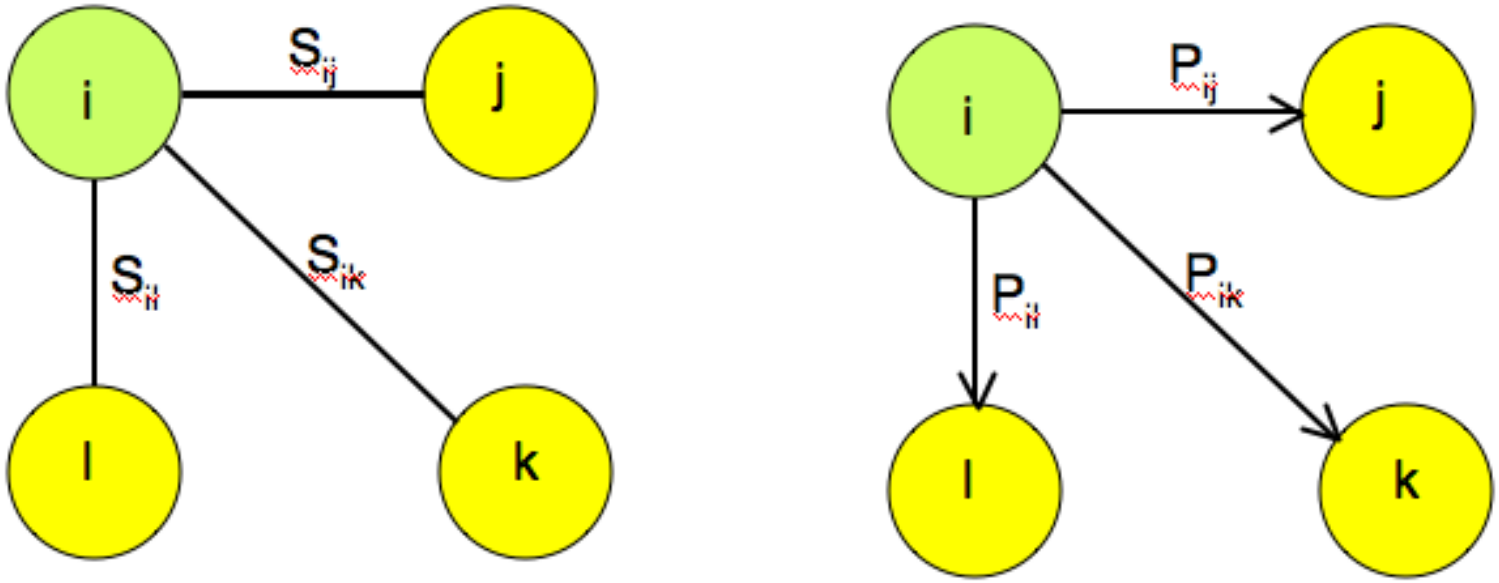

- Define

$$P_{ij} = \frac{S_{ij}}{D_i} \quad \text{for all } i, j \in V$$

- in matrix notation $P = D^{-1}S$ where $P = [P_{ij}]$, $D = \text{diag}(D_1, \ldots D_n)$
- P defines a **random walk** over the graph nodes V

Grouping from the random walks point of view

- Idea: group nodes together if they transition in the same way to other clusters

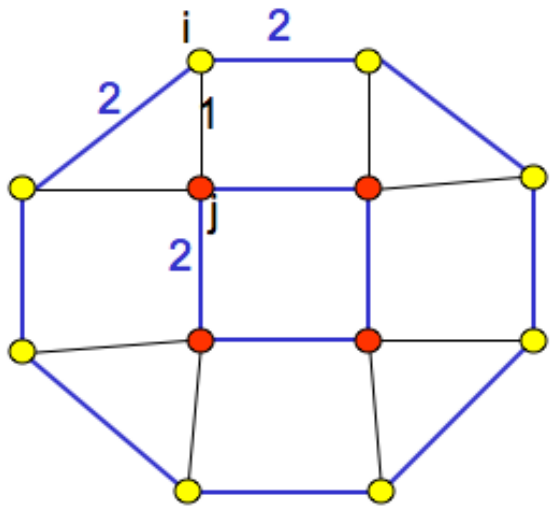

$$P_{i,red} = Pr[i \rightarrow red \,|\, i] = \sum_{j \in red} P_{ij}$$

i	$P_{i,red}$	$P_{i,yellow}$
●	1/5	4/5
●	1/5	4/5
●	1/5	4/5
●	1/5	4/5
●	1/5	4/5
●	1/5	4/5
●	1/5	4/5
●	1/5	4/5
●	2/3	1/3
●	2/3	1/3
●	2/3	1/3
●	2/3	1/3

... is the same as grouping by embedding

- **embedding** of V = mapping from V into $\mathbb{R}^d$
- Wanted: similar points embedded near each other
  ideally, points in the same cluster mapped to the same point in $\mathbb{R}^d$

Another look at $P_{i,C}$

a piecewise constant function

$P_{\cdot,red} \equiv f^{red}$

$P_{\cdot,yel} \equiv f^{yel}$

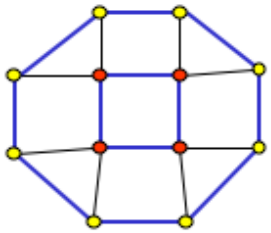

i	$P_{i,red}$	$P_{i,yellow}$
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	2/3	1/3
	2/3	1/3
	2/3	1/3
	2/3	1/3

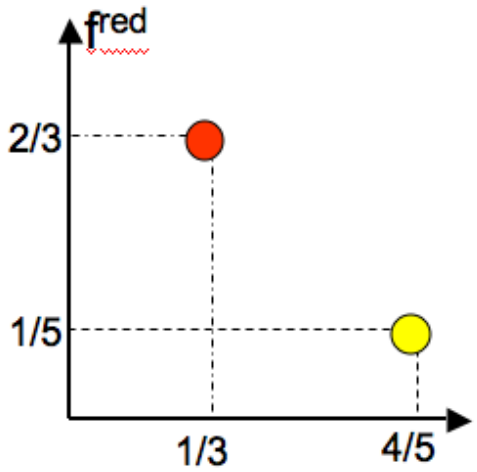

Some questions

- Not all graphs embed perfectly

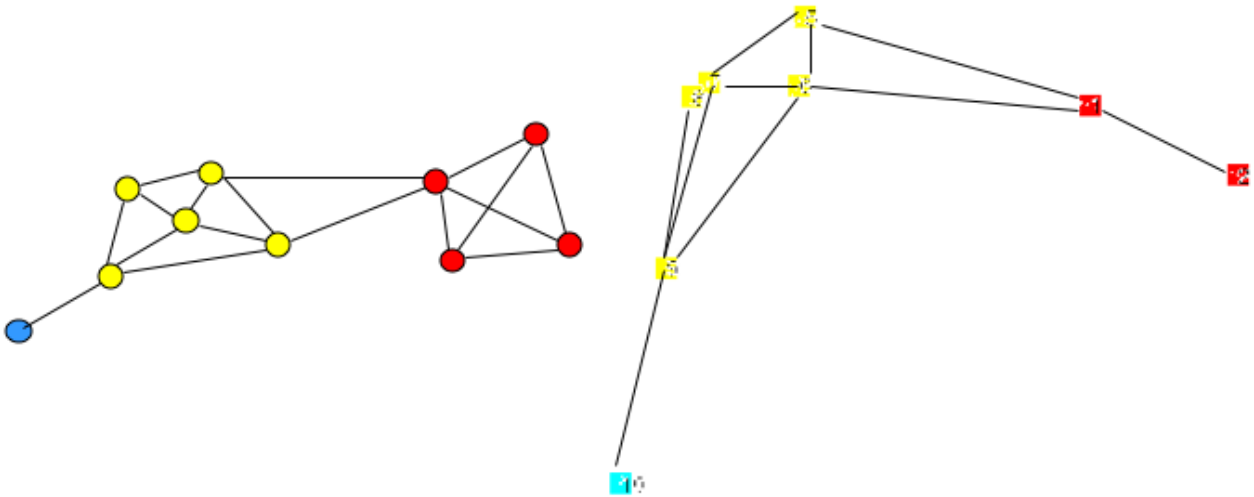

- How many dimensions do we need?
- Nice, but we need to know the clusters in advance. . .

Lumpability

- A vector v is **piecewise constant** w.r.t a clustering Δ iff $v_i = v_j$ whenever i, j in same $C \in \Delta$

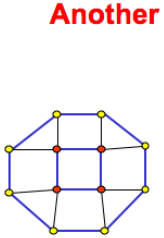

Another look at $P_{i,C}$

$P_{\cdot,red} \equiv f^{red}$

$P_{\cdot,yel} \equiv f^{yel}$

a piecewise constant function

i	$P_{i,red}$	$P_{i,yellow}$
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	1/5	4/5
	2/3	1/3
	2/3	1/3
	2/3	1/3
	2/3	1/3

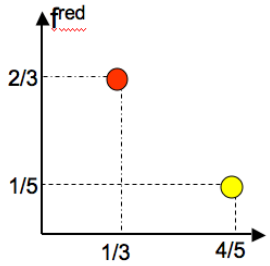

- **Theorem [Lumpability]**[Meila&Shi 2001] Let S be a similarity matrix and Δ a clustering with K clusters. Then P has K **piecewise constant** eigenvectors w.r.t Δ **iff**

$$\sum_{j \in C'} P_{ij} \;=\; R_{CC'} \quad \text{whenver } i \in C, \text{ for all } C, C' \in \Delta$$

The spectral mapping

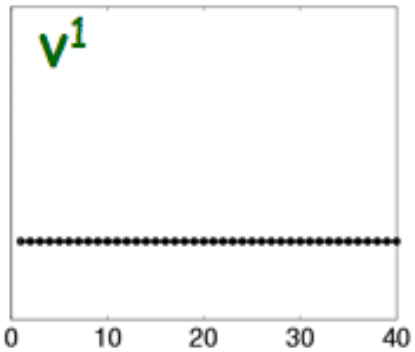

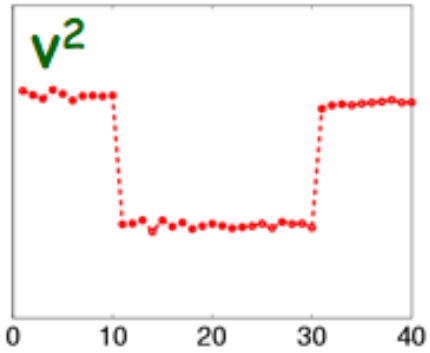

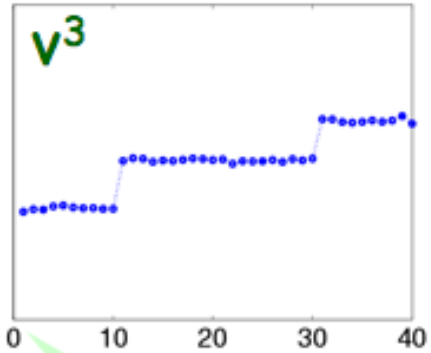

The **spectral mapping**: Data as elements of v^2, v^3

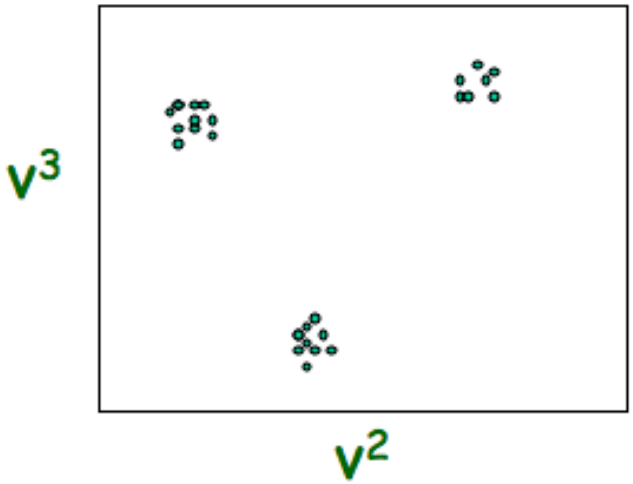

These eigenvectors are called **piecewise constant (PC)**

Spectral clustering in a nutshell

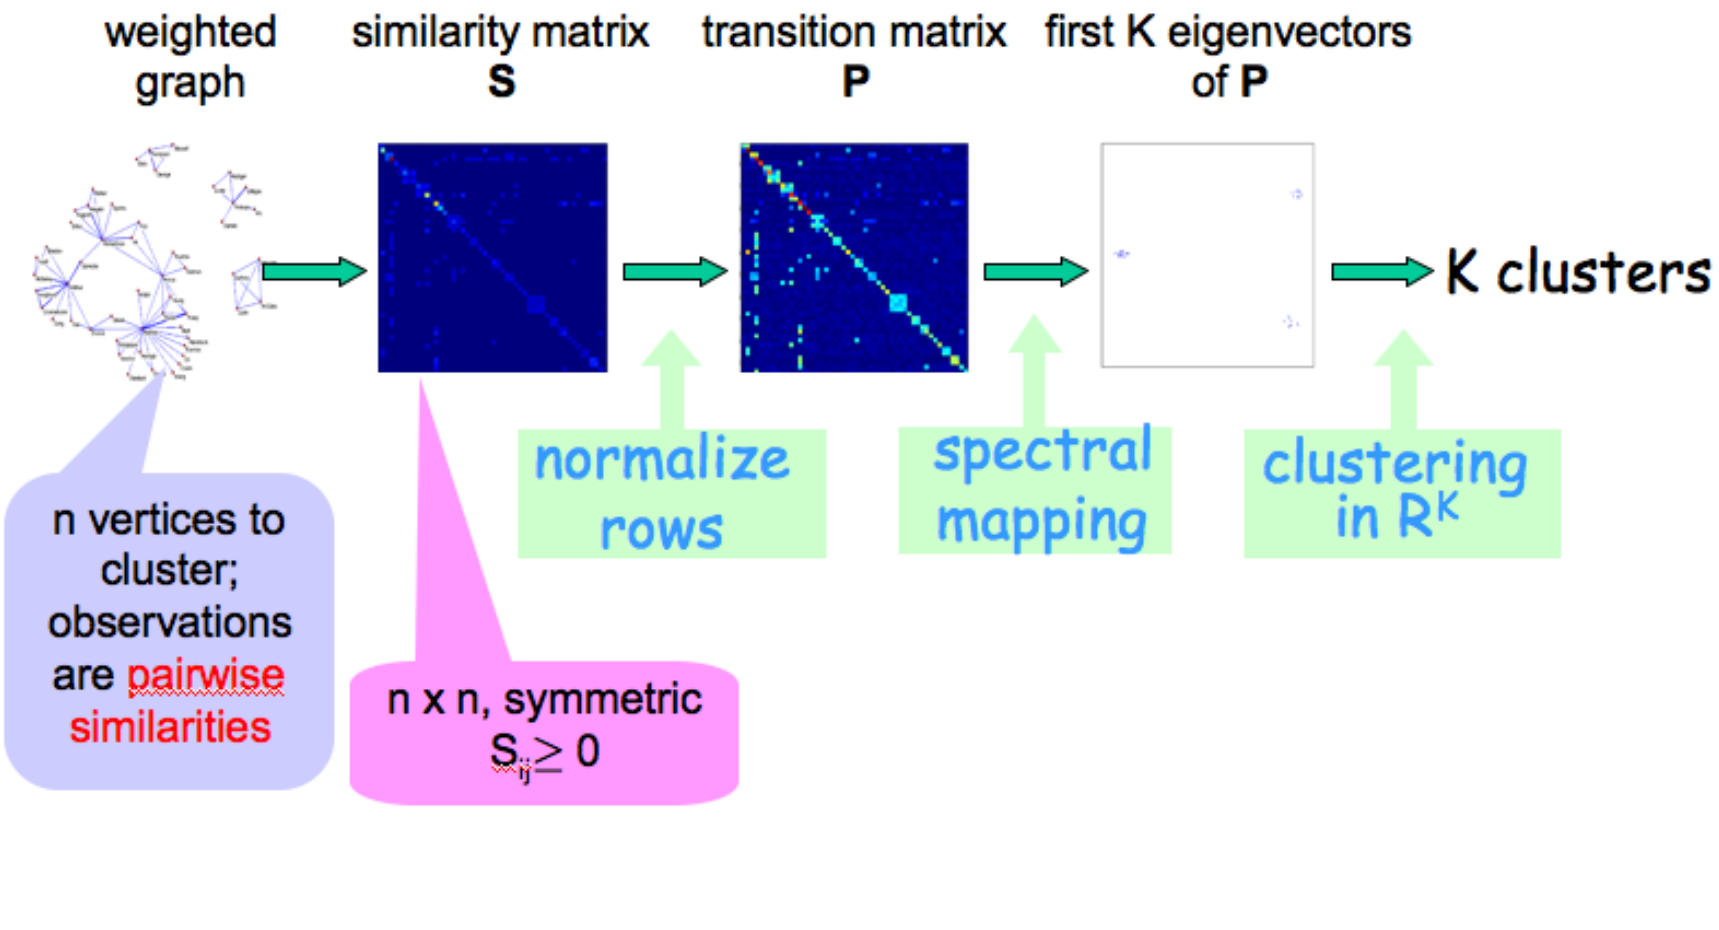

Spectral clustering

An algorithm based on [?] and [?].

Spectral Clustering Algorithm

Input Similarity matrix S, number of clusters K

1. Transform S: Set $D_i = \sum_{j=1}^{n} S_{ij}$, $j = 1 : n$ the **node degrees**.
   Form the **transition matrix** $P = [P_{ij}]_{ij=1}^{n}$ with

$$P_{ij} \leftarrow S_{ij}/D_i, \text{ for } i, j = 1 : n$$

2. Compute the largest K eigenvalues $\lambda_1 = 1 \geq \lambda_2 \geq \ldots \geq \lambda_K$ and eigenvectors $\mathrm{v}_1, \ldots \mathrm{v}_K$ of P.
3. Embed the data in principal subspace Let $V = [\,\mathrm{v}_2\ \mathrm{v}_3\ \ldots\ \mathrm{v}_K\,] \in \mathbb{R}^{n \times K}$, $\mathrm{x}_i \leftarrow$ i-th row of V.
4. (orthogonal initialization) Find K initial centers by
   1. take μ_1 randomly from $\mathrm{x}_1, \ldots \mathrm{x}_n$
   2. for $k = 2, \ldots K$ set $\mu_k = \mathrm{argmin}_{\mathrm{x}_i} \max_{k' < k} \mu_{k'}^T \mathrm{x}_i$.
5. Run the K-means algorithm on the "data" $\mathrm{x}_{1:n}$ starting from the centers $\mu_{1:K}$.

Properties of spectral clustering

- Arbitrary cluster shapes (main advantage)
- Elegant mathematically
- Practical up to medium sized problems
  - Running time (by Lanczos algorithm) $\mathcal{O}(nk)$/iteration.
- Works well when K known, not too large
  estimating K [**?**]
- Depend heavily on the similarity function (main problem)
  learning the similarities [**?**],[**?**],[**?**],[**?**]
- Outliers become separate clusters (user must adjust K accordingly!)
- Very popular, many variants which aim to improve on the above
  Diffusion maps [**?**]: normalize the eigenvectors $\lambda_k^t v^k$
- Practical fix, when K large: only compute a fixed number of eigenvectors $d < K$. This avoids the effects of noise in lower ranked eigenvectors

Affinity propagation

- **Idea** Each item $i \in \mathcal{D}$ finds an **exemplar** item $k \in \mathcal{D}$ to "represent" it
- Affinity Propagation is to spectral clustering what Mean Shift is to K-means
- number of exemplars not fixed in advance
- quantities of interest
  - similarities s_{ij}, $i \neq j$ (given)
  - **availability** a_{ik} of k for i = how much support there is from other items for k to be an exemplar
  - **responsibility** r_{ik} that measures how fit is k to represent i, as compared to other possible candidates k'.
  - diagonal elements s_{ii} represent **self-similarities**
    - larger s_{ii} $\Rightarrow$ more likely i will become an exemplar $\Rightarrow$ more clusters

Affinity Propagation

Affinity Propagation Algorithm [?]

Input Similarity matrix $S = [s_{ik}]_{ik=1}^{n}$, parameter $\lambda = 0.5$

Iterate the following steps until convergence

1. $a_{ik} \leftarrow 0$ for $i, k = 1 : n$
2. for all i
   1. Find the best exemplar for i: $s^* \leftarrow \max_k (s_{ik} + a_{ik})$, $A_i^* \leftarrow \operatorname{argmax}_k (s_{ik} + a_{ik})$ (can be a set of items)
   2. for all k update responsibilities

$$r_{ik} \leftarrow \begin{cases} s_{ik} - s^*, & \text{if } k \notin A_i^* \\ s_{ik} - \max_{k' \notin A_i^*} (s_{ik} + a_{ik}) & \text{otherwise} \end{cases}$$

3. for all k update availabilities
   1. $a_{kk} \leftarrow \sum_{i \neq k} [r_{ik}]_+$ where $[r_{ik}]_+ = r_{ik}$ if $r_{ik} > 0$ and 0 otherwise.
   2. for all i, $a_{ik} \leftarrow \min\{0, r_{kk} + \sum_{i' \neq i,k} [r_{i'k}]_+\}$
4. Assign an exemplar to i by $k(i) \leftarrow \operatorname{argmax}_{k'} (r_{ik'} + a_{ik'})$